



OBJECT MANAGEMENT GROUP®

Extensible and Dynamic Topic Types for DDS

Version 1.2

OMG Document Number: formal/2017-08-01

Date: August 2017

Standard document URL: <http://www.omg.org/spec/DDS-XTypes/1.2/>

Normative Machine Consumable File(s):

http://www.omg.org/spec/DDS-XTypes/20170301/dds-xtypes_model.xmi

http://www.omg.org/spec/DDS_XTypes/20170301/dds-xtypes_type_definition.xsd

http://www.omg.org/spec/DDS_XTypes/20170301/dds-xtypes_type_definition_nonamespace.xsd

http://www.omg.org/spec/DDS-XTypes/20170301/dds-xtypes_typeobject.idl

http://www.omg.org/spec/DDS_XTypes/20170301/dds-xtypes_discovery.idl

Copyright ©2017, Object Management Group, Inc.

Copyright ©2008-2017, PrismTech Group Ltd.

Copyright ©2008-2017, Real-Time Innovations, Inc.

Copyright ©2008-2017, Twin Oaks Computing, Inc.

Copyright ©2008-2017, Object Computing, Inc.

USE OF SPECIFICATION - TERMS, CONDITIONS & NOTICES

The material in this document details an Object Management Group specification in accordance with the terms, conditions and notices set forth below. This document does not represent a commitment to implement any portion of this specification in any company's products. The information contained in this document is subject to change without notice.

LICENSES

The companies listed above have granted to the Object Management Group, Inc. (OMG) a nonexclusive, royalty-free, paid up, worldwide license to copy and distribute this document and to modify this document and distribute copies of the modified version. Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having used the specification set forth herein or having conformed any computer software to the specification.

Subject to all of the terms and conditions below, the owners of the copyright in this specification hereby grant you a fully-paid up, non-exclusive, nontransferable, perpetual, worldwide license (without the right to sublicense), to use this specification to create and distribute software and special purpose specifications that are based upon this specification, and to use, copy, and distribute this specification as provided under the Copyright Act; provided that: (1) both the copyright notice identified above and this permission notice appear on any copies of this specification; (2) the use of the specifications is for informational purposes and will not be copied or posted on any network computer or broadcast in any media and will not be otherwise resold or transferred for commercial purposes; and (3) no modifications are made to this specification. This limited permission automatically terminates without notice if you breach any of these terms or conditions. Upon termination, you will destroy immediately any copies of the specifications in your possession or control.

PATENTS

The attention of adopters is directed to the possibility that compliance with or adoption of OMG specifications may require use of an invention covered by patent rights. OMG shall not be responsible for identifying patents for which a license may be required by any OMG specification, or for conducting legal inquiries into the legal validity or scope of those patents that are brought to its attention. OMG specifications are prospective and advisory only. Prospective users are responsible for protecting themselves against liability for infringement of patents.

GENERAL USE RESTRICTIONS

Any unauthorized use of this specification may violate copyright laws, trademark laws, and communications regulations and statutes. This document contains information which is protected by copyright. All Rights Reserved. No part of this work covered by copyright herein may be reproduced or used in any form or by any means--graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems--without permission of the copyright owner.

DISCLAIMER OF WARRANTY

WHILE THIS PUBLICATION IS BELIEVED TO BE ACCURATE, IT IS PROVIDED "AS IS" AND MAY CONTAIN ERRORS OR MISPRINTS. THE OBJECT MANAGEMENT GROUP AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS PUBLICATION, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, IMPLIED WARRANTY OF MERCHANTABILITY OR WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE OR USE. IN NO EVENT SHALL THE OBJECT MANAGEMENT GROUP OR ANY OF THE COMPANIES LISTED ABOVE BE LIABLE FOR ERRORS CONTAINED HEREIN OR FOR DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, RELIANCE OR COVER DAMAGES, INCLUDING LOSS OF PROFITS, REVENUE, DATA OR USE, INCURRED BY ANY USER OR ANY THIRD PARTY IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The entire risk as to the quality and performance of software developed using this specification is borne by you. This disclaimer of warranty constitutes an essential part of the license granted to you to use this specification.

RESTRICTED RIGHTS LEGEND

Use, duplication or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c) (1) (ii) of The Rights in Technical Data and Computer Software Clause at DFARS 252.227-7013 or in subparagraph (c)(1) and (2) of the Commercial Computer Software - Restricted Rights clauses at 48 C.F.R. 52.227-19 or as specified in 48 C.F.R. 227-7202-2 of the DoD F.A.R. Supplement and its successors, or as specified in 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors, as applicable. The specification copyright owners are as indicated above and may be contacted through the Object Management Group, 109 Highland Avenue, Needham, MA 02494, U.S.A.

TRADEMARKS

C®, CORBA®, CORBA logos®, FIBO®, Financial Industry Business Ontology®, FINANCIAL INSTRUMENT GLOBAL IDENTIFIER®, IIOP®, IMM®, Model Driven Architecture®, MDA®, Object Management Group®, OMG®, OMG Logo®, SoaML®, SOAML®, SysML®, UAF®, Unified Modeling Language®, UML®, UML Cube logo®, VSIPL®, and XMI® are registered trademarks of the Object Management Group, Inc.

For a complete list of trademarks, see: http://www.omg.org/legal/tm_list.htm. All other products or company names mentioned are used for identification purposes only, and may be trademarks of their respective owners.

COMPLIANCE

The copyright holders listed above acknowledge that the Object Management Group (acting itself or through its designees) is and shall at all times be the sole entity that may authorize developers, suppliers and sellers of computer software to use certification marks, trademarks or other special designations to indicate compliance with these materials.

Software developed under the terms of this license may claim compliance or conformance with this specification if and only if the software compliance is of a nature fully matching the applicable compliance points as stated in the specification. Software developed only partially matching the applicable compliance points may claim only that the software was based on this specification, but may not claim compliance or conformance with this specification. In the event that testing suites are implemented or approved by Object Management Group, Inc., software developed using this specification may claim compliance or conformance with the specification only if the software satisfactorily completes the testing suites.

OMG's Issue Reporting

All OMG specifications are subject to continuous review and improvement. As part of this process we encourage readers to report any ambiguities, inconsistencies, or inaccuracies they may find by completing the Issue Reporting Form.

Table of Contents

Table of Contents	v
Tables	x
Figures	xiii
About the Object Management Group	xv
OMG	xv
OMG Specifications	xv
OMG Modeling Specifications	xv
OMG Middleware Specifications	xv
Platform Specific Model and Interface Specifications	xvi
Typographical Conventions	xvi
Issues	xvi
1. Scope	1
2. Conformance Criteria	3
2.1 Programming Interface Conformance	3
2.2 Network Interoperability Conformance	4
2.2.1 Minimal Network Interoperability Profile	4
2.2.2 Basic Network Interoperability Profile	4
2.3 Optional XTYPES 1.1 Interoperability Profile	4
2.4 Optional XML Data Representation Profile	5
3. Normative References	7
4. Terms and Definitions	9
5. Symbols	11
6. Additional Information	13
6.1 Data Distribution Service for Real-Time Systems (DDS)	13
6.2 Acknowledgments	15
7. Extensible and Dynamic Topic Types for DDS	17
7.1 Overview	17
7.2 Type System	19
7.2.1 Background (Non-Normative)	19
7.2.1.1 Type Evolution Example	20
7.2.1.2 Type Inheritance Example	21

7.2.1.3	Sparse Types Example.....	22
7.2.2	Type System Model	23
7.2.2.1	Namespaces.....	23
7.2.2.2	Primitive Types.....	24
7.2.2.3	String Types.....	30
7.2.2.4	Constructed Types	31
7.2.2.5	Nested Types.....	48
7.2.2.6	Annotations.....	48
7.2.2.7	Try Construct behavior	49
7.2.3	Type Extensibility and Mutability.....	52
7.2.4	Type Compatibility	53
7.2.4.1	Constructing objects of one type from objects of another type	54
7.2.4.2	Concept of Delimited Types	55
7.2.4.3	Strong Assignability.....	55
7.2.4.4	Assignability Rules	55
7.3	Type Representation.....	66
7.3.1	IDL Type Representation.....	68
7.3.1.1	IDL Compatibility.....	68
7.3.1.2	Annotation Language.....	70
7.3.1.3	Constants and Expressions.....	83
7.3.1.4	Primitive Types.....	83
7.3.1.5	Alias Types	84
7.3.1.6	Array and Sequence Types	84
7.3.1.7	String Types.....	84
7.3.1.8	Enumerated Types	84
7.3.1.9	Map Types	84
7.3.1.10	Structure Types	84
7.3.1.11	Union Types.....	84
7.3.2	XML Type Representation.....	85
7.3.2.1	Type Representation Management	85
7.3.2.2	Basic Types.....	86
7.3.2.3	String Types.....	87
7.3.2.4	Collection Types	88

7.3.2.5	Aggregated Types	90
7.3.2.6	Aliases	93
7.3.2.7	Enumerated Types	93
7.3.2.8	Modules.....	93
7.3.2.9	Annotations	94
7.3.3	XSD Type Representation.....	94
7.3.3.1	Annotations	94
7.3.3.2	Structures	96
7.3.3.3	Nested Types.....	97
7.3.3.4	Maps.....	97
7.3.4	Representing Types with TypeIdentifier and TypeObject	98
7.3.4.1	Plain Types.....	98
7.3.4.2	Type Identifier	98
7.3.4.3	Complete TypeObject	100
7.3.4.4	Minimal TypeObject.....	100
7.3.4.5	TypeObject serialization	101
7.3.4.6	Classification of TypeIdentifiers.....	102
7.3.4.7	Type Equivalence.....	103
7.3.4.8	Types with mutual dependencies on other types	104
7.3.4.9	Computation of Type identifiers for types with mutual dependencies	106
7.4	Data Representation.....	109
7.4.1	Extended CDR Representation (encoding version 1)	111
7.4.1.1	PLAIN_CDR Representation.....	111
7.4.1.2	Parameterized CDR Representation	114
7.4.2	Extended CDR Representation (encoding version 2)	119
7.4.3	Extended CDR encoding virtual machine	120
7.4.3.1	Encoding version and format.....	120
7.4.3.2	XCDR Stream State	120
7.4.3.3	Type and Byte transformations.....	123
7.4.3.4	Functions related to data types and objects.....	125
7.4.3.5	Encoding (serialization) rules	128
7.4.4	XML Data Representation	141
7.4.4.1	Valid XML Data Representation	142

7.4.4.2	Well-formed XML Data Representation	142
7.5	Language Binding.....	143
7.5.1	Plain Language Binding.....	144
7.5.1.1	Primitive Types.....	145
7.5.1.2	Annotations and Built-in Annotations	147
7.5.1.3	Map Types	158
7.5.1.4	Structure and Union Types	166
7.5.2	Dynamic Language Binding.....	167
7.5.2.1	UML-to-IDL Mapping Rules.....	168
7.5.2.2	DynamicTypeBuilderFactory	170
7.5.2.3	AnnotationDescriptor.....	176
7.5.2.4	TypeDescriptor	178
7.5.2.5	MemberId.....	181
7.5.2.6	DynamicTypeMember	181
7.5.2.7	MemberDescriptor	183
7.5.2.8	DynamicType.....	186
7.5.2.9	DynamicTypeBuilder.....	189
7.5.2.10	DynamicDataFactory	192
7.5.2.11	DynamicData	193
7.6	Use of the Type System by DDS.....	199
7.6.1	Topic Model.....	199
7.6.2	Discovery and Endpoint Matching.....	200
7.6.2.1	Data Representation QoS Policy.....	200
7.6.2.2	Discovery Built-in Topics.....	206
7.6.2.3	Built-in TypeLookup service	209
7.6.2.4	Type Consistency Enforcement QoS Policy	216
7.6.3	Local API Extensions.....	219
7.6.3.1	Operation: DomainParticipant::create_topic.....	219
7.6.3.2	Operation: DomainParticipant::lookup_topicdescription..	219
7.6.4	Built-in Types.....	220
7.6.4.1	String.....	220
7.6.4.2	KeyedString	221

7.6.4.3	Bytes	221
7.6.4.4	KeyedBytes	221
7.6.5	Use of Dynamic Data and Dynamic Type	221
7.6.5.1	Type Support.....	221
7.6.5.2	DynamicDataWriter and DynamicDataReader	224
7.6.6	DCPS Queries and Filters	224
7.6.6.1	Member Names.....	224
7.6.6.2	Optional Type Members	225
7.6.6.3	Grammar Extensions.....	225
7.6.7	Interoperability of Keyed Topics	225
8.	Changes or Extensions Required to Adopted OMG Specifications	227
8.1	Extensions.....	227
8.1.1	DDS.....	227
8.2	Changes	227
Annex A:	XML Type Representation Schema	229
Annex B:	Representing Types with TypeObject	247
Annex C:	Dynamic Language Binding.....	277
Annex D:	DDS Built-in Topic Data Types.....	291
Annex E:	Built-in Types	301
Annex F:	Characterizing Legacy DDS Implementations	307
F.1	Type System.....	307
F.2	Type Representation.....	307
F.3	Data Representation	308
F.4	Language Binding	308

Tables

Table 1 – Type-related concerns addressed by this specification.....	14
Table 2 – Main features and mechanisms provided by this Specification to address type-related concerns.....	14
Table 3 – Primitive Types.....	26
Table 4 – Enumerated types.....	33
Table 5 – Bitmask types.....	34
Table 6 – Alias types.....	36
Table 7 – Collection Types.....	38
Table 8 – Aggregated Types.....	40
Table 9 – Default values for non-optional members.....	44
Table 10 – TryConstruct examples.....	50
Table 11 – TryConstruct behavior kinds.....	51
Table 12 – Meaning of marking types as appendable.....	53
Table 13 – Type assignability example.....	54
Table 14 – Definition of the <i>is-assignable-from</i> relationship for alias types.....	56
Table 15 – Definition of the <i>is-assignable-from</i> relationship for primitive types.....	56
Table 16 – Definition of the <i>is-assignable-from</i> relationship for string types.....	57
Table 17 – Definition of the <i>is-assignable-from</i> relationship for collection types.....	58
Table 18 – Definition of the <i>is-assignable-from</i> relationship for alias, bitmask, and enumerated types.....	60
Table 19 – Definition of the <i>is-assignable-from</i> relationship for aggregated types.....	60
Table 20 – Alternative Type Representations.....	67
Table 21 – IDL Built-in Annotations Usage.....	78
Table 22 – Syntax for declaring an annotation type.....	80
Table 23 – Syntax for members of annotation types.....	80
Table 24 – Syntax for applying annotations.....	81
Table 25 – IDL primitive type mapping.....	83
Table 26 – Primitive and string type names in the XML Type Representation.....	87
Table 27– XSD annotation example.....	95
Table 28 – XSD structure inheritance example.....	97
Table 29 – Formats and interpretation of the TypeIdentifier.....	98

Table 31 – Serialization of primitive types in version 1 encoding	112
Table 32 – Serialization of enumeration types	113
Table 33 – Serialization of bitmask types	113
Table 34 – Reserved parameter ID values	116
Table 35 – Serialization format to use	120
Table 36 – State variables and constants in the XCDR stream model.....	121
Table 37 – Stream operations in the XCDR stream model.....	122
Table 38 – Type and Byte transformations used in the serialization virtual machine	124
Table 39 – Functions operating on objects and types	125
Table 40 – Symbols and notation used in the serialization virtual machine.....	129
Table 41 – Kinds of Language Bindings	144
Table 42 – Plain Language Binding for Primitive Types in C	145
Table 43 – Plain Language Binding for Primitive Types in C++	146
Table 44 – Bit mask integer equivalents.....	148
Table 45 – Configurable behaviors of the copy function when destination is not NULL	150
Table 46 – Behavior of assignment operator	154
Table 47 – Operations for map<KeyType, ElementType>	158
Table 48 – DynamicTypeBuilderFactory properties and operations	170
Table 49 – AnnotationDescriptor properties and operations	176
Table 50 – TypeDescriptor properties and operations	179
Table 51 – DynamicMember behavior	181
Table 52 – DynamicTypeMember properties and operations.....	182
Table 53 – MemberDescriptor properties and operations.....	183
Table 54 – DynamicType properties and operations	187
Table 55 – DynamicType::member_by_name behavior.....	188
Table 56 – DynamicTypeBuilder properties and operations	190
Table 57 – DynamicDataFactory properties and operations.....	192
Table 58 – DynamicData properties and operations.....	194
Table 59 – Compatibility matrix for the DataRepresentationQosPolicy	203
Table 60 – RTPS encapsulation identifier	204
Table 61 – Built-in Endpoints added by the XTYPES specification.....	211
Table 62 – Mapping of the built-in endpoints added by this specification to the availableBuiltinEndpoints.....	214

Table 63 – New TypeSupport operations	222
Table 64 – New FooTypeSupport operations	222
Table 65 – DynamicTypeSupport properties and operations	223

Figures

Figure 1 – Packages	1
Figure 2 – Relationships between Type System, Type Representation, Language Binding, and Data Representation	17
Figure 3 – Example Type Representation, Language Binding, and Data Representation.....	19
Figure 4 – Type System Model.....	23
Figure 5 – Namespaces	24
Figure 6 – Primitive Types: Integral Types	25
Figure 7 – Primitive Types: Floating Point Types.....	25
Figure 8 – Primitive Types: Booleans, Bytes, and Characters	26
Figure 9 - String Types	30
Figure 10 – Constructed Types	31
Figure 11 – Enumerated Types	32
Figure 12 – Enumeration Types.....	33
Figure 13 – Bitmask Types	34
Figure 14 – Alias Types.....	36
Figure 15 – Collection Types.....	37
Figure 16 – Aggregated Types.....	40
Figure 17 – Structure Types.....	41
Figure 18 – Union Types	42
Figure 19 – Annotation Types	49
Figure 20 – Type Representation.....	66
Figure 21 – Directed graph, Strongly Connected Components, and Kernel DAG.....	106
Figure 22 – Dependency graph derived from a set of type definitions.....	106
Figure 23 – Data Representation—conceptual model	110
Figure 24 – Usage of PID_EXTENDED within the CDR Buffer	118
Figure 25 – Language Bindings—conceptual model.....	143
Figure 26 – Dynamic Data and Dynamic Type	168
Figure 27 – Annotation Descriptor	176
Figure 28 – Type Descriptor.....	178
Figure 29 – Dynamic Type Members	182
Figure 30 – Dynamic Type	187

Figure 31 – Dynamic Data and Dynamic Data Factory.....	194
Figure 32 – Dynamic Type Support.....	223

Preface

About the Object Management Group

OMG

Founded in 1989, the Object Management Group, Inc. (OMG) is an open membership, not-for-profit computer industry standards consortium that produces and maintains computer industry specifications for interoperable, portable and reusable enterprise applications in distributed, heterogeneous environments. Membership includes Information Technology vendors, end users, government agencies, and academia.

OMG member companies write, adopt, and maintain its specifications following a mature, open process. OMG's specifications implement the Model Driven Architecture® (MDA®), maximizing ROI through a full-lifecycle approach to enterprise integration that covers multiple operating systems, programming languages, middleware and networking infrastructures, and software development environments. OMG's specifications include: UML® (Unified Modeling Language™); CORBA® (Common Object Request Broker Architecture); CWM™ (Common Warehouse Metamodel); and industry-specific standards for dozens of vertical markets.

More information on the OMG is available at <http://www.omg.org/>.

OMG Specifications

As noted, OMG specifications address middleware, modeling, and vertical domain frameworks. A catalog of all OMG Specifications Catalog is available from the OMG website at:

http://www.omg.org/technology/documents/spec_catalog.htm

Specifications within the Catalog are organized by the following categories:

OMG Modeling Specifications

- UML
- MOF
- XMI
- CWM
- OMG SysML™
- Other Profile specifications

OMG Middleware Specifications

- CORBA/IIOP
- DDS and the DDS Interoperability Protocol, RTPS
- IDL/Language Mappings
- Specialized CORBA specifications

- CORBA Component Model (CCM)

Platform Specific Model and Interface Specifications

- CORBA services
- CORBA facilities
- OMG Domain specifications
- OMG Embedded Intelligence specifications
- OMG Security specifications

All of the OMG's formal specifications may be downloaded without charge from our website. (Products implementing OMG specifications are available from individual suppliers.) Copies of specifications, available in PostScript and PDF format, may be obtained from the Specifications Catalog cited above or by contacting the Object Management Group, Inc. at:

OMG Headquarters
109 Highland Avenue
Needham, MA 02494, USA
Tel: +1-781-444-0404
Fax: +1-781-444-0320
Email: pubs@omg.org

Certain OMG specifications are also available as ISO standards. Please consult <http://www.iso.org>.

Typographical Conventions

The type styles shown below are used in this document to distinguish programming statements from ordinary English. However, these conventions are not used in tables or section headings where no distinction is necessary.

Times/Times New Roman - 10 pt.: Standard body text

Helvetica/Arial - 10 pt. Bold: OMG Interface Definition Language (OMG IDL) and syntax elements.

Courier - 10 pt. Bold: Programming language elements.

Helvetica/Arial - 10 pt: Exceptions

Terms that appear in italics are defined in the glossary. Italic text also represents the name of a document, specification, or other publication.

Issues

The reader is encouraged to report any technical or editing issues/problems with this specification to <http://issues.omg.org/issues/create-new-issue>.

1. Scope

The Specification addresses four related concerns summarized in Figure 1 below.

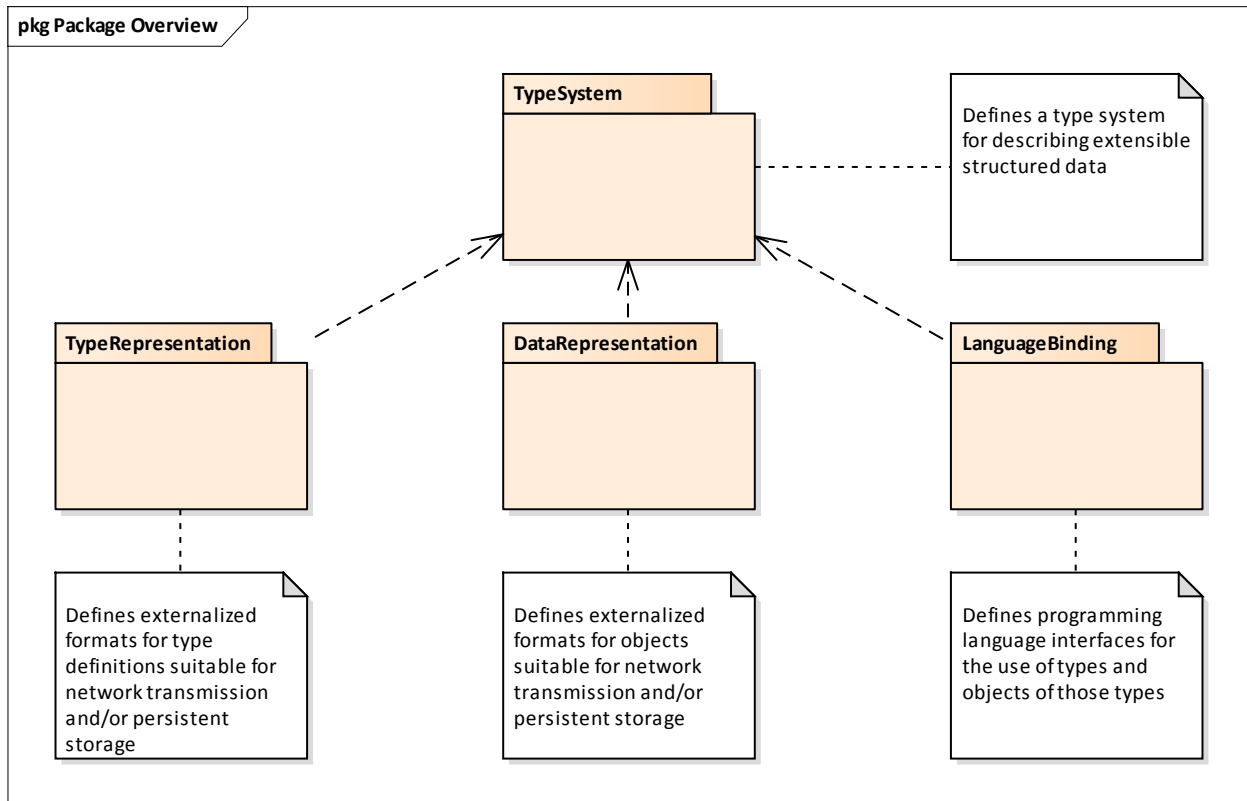


Figure 1 – Packages

The specification addresses four related concerns: the type system, the representation of types, the representation of data, and the language bindings used to access types and data. Each of these concerns is modeled as a collection of classes belonging to a corresponding package.

This specification provides the following additional facilities to DDS [DDS] implementations and users:

- **Type System.** The specification defines a model of the data types that can be used for DDS Topics. The type system is formally defined using UML. The Type System is defined in Clause 7.2 and its sub clauses. The structural model of this system is defined in the Type System Model in Clause 7.2.2. The framework under which types can be modified over time is summarized in Clause 7.2.3, “Type Extensibility and Mutability.” The concrete rules under which the concepts from 7.2.2 and 7.2.3 come together to define compatibility in the face of such modifications are defined in Clause 7.2.4, “Type Compatibility.”
- **Type Representations.** The specification defines the ways in which types described by the Type System may be externalized such that they can be stored in a file or communicated over a network. The specification adds additional Type Representations

beyond the one (IDL [IDL41]) already implied by the DDS specification. Several Type Representations are specified in the sub clauses of Clause 7.3. These include IDL (7.3.1), XML (7.3.2), XML Schema (XSD) (7.3.3), and TypeObject (7.3.4).

- **Data Representation.** The specification defines multiple ways in which objects of the types defined by the Type System may be externalized such that they can be stored in a file or communicated over a network. (This is also commonly referred as “data serialization” or “data marshaling.”) The specification extends and generalizes the mechanisms already defined by the DDS Interoperability specification [RTPS]. The specification includes Data Representations that support data type evolution, that is, allow a data type to change in certain well-defined ways without breaking communication. Two Data Representations are specified in the sub clauses of Clause 7.4. These are Extended CDR (7.4.1, 7.4.2, and 7.4.3) and XML (7.4.4).
- **Language Binding.** The specification defines multiple ways in which applications can access the state of objects defined by the Type System. The specification extends and generalizes the mechanism currently implied by the DDS specification (“Plain Language Binding”) and adds a Dynamic Language Binding that allows application to access data without compile-time knowledge of its type. The specification also defines an API to define and manipulate data types programmatically. Two Language Bindings are specified in the sub clauses of Clause 7.5. These are the Plain Language Binding and the Dynamic Language Binding.

2. Conformance Criteria

This specification recognizes two areas of conformance: (1) conformance with respect to programming interfaces—that is, at the level of the DDS API—and (2) conformance with respect to network interoperability—that is, at the level of the RTPS protocol.

Additionally, it defines two optional profiles: XTYPES 1.1 Interoperability and XML Data Representation.

There are three conformance levels:

- *Minimal* conformance with XTYPES version 1.2 requires conformance to the Programming Interface and the Minimal Network Interoperability Profile.
- *Basic* conformance with XTYPES version 1.2 requires conformance to the Programming Interface and the Basic Network Interoperability Profile.
- *Complete* conformance with XTYPES version 1.2 requires Basic conformance as well as conformance to the two optional profiles.

2.1 Programming Interface Conformance

This specification extends the *Data Distribution Service for Real-Time Systems* specification [DDS] with an additional optional conformance profile: the “Extensible and Dynamic Types Profile.” Conformance to this specification with respect to programming interfaces shall be equivalent to conformance to the DDS specification with respect to at least the existing Minimum Profile and the new Extensible and Dynamic Types Profile. Implementations may conform to additional DDS profiles.

The new Extensible and Dynamic Types profile of DDS shall consist of the following clauses of this specification:

- “Extensible and Dynamic Topic Types for DDS” (Clause 7) up to and including “Type Representation” (Clause 7.3)
- “Language Binding” (Clause 7.5)
- “Use of the Type System by DDS” (Clause 7.6) excluding “Interoperability of Keyed Topics” (Clause 7.6.7)
- Annex B: Representing Types with TypeObject
- Annex C: Dynamic Language Binding
- Annex E: Built-in Types

2.2 Network Interoperability Conformance

There are two Network Interoperability conformance profiles. An implementation may claim conformance to the Minimal profile or to the Basic profile, which extends the Minimal.

Regardless of profile, conformance with respect to network interoperability requires conformance to the *Real-Time Publish-Subscribe Wire Protocol* specification [RTPS].

2.2.1 Minimal Network Interoperability Profile

Conformance with the Minimal Network Interoperability profile requires conformance with the following clauses of this specification:

- “Representing Types with TypeIdentifier and TypeObject” (Clause 7.3.4)
- From “Use of the Type System by DDS” (Clause 7.6)
 - “Topic Model” (Clause 7.6.1)
 - “Discovery and Endpoint Matching” (Clause 7.6.2) excluding “Built-in TypeLookup service” (Clause 7.6.2.3)
 - Clause 7.6.2.1.1 “DataRepresentationQosPolicy: Conceptual Model”, with support limited to version 2 encoding.
 - “Interoperability of Keyed Topics” (Clause 7.6.7)
- “Extended CDR Representation (encoding version 2)” (Clause 7.4.2)
- “Extended CDR encoding virtual machine” (Clause 7.4.3)
- Annex B: Representing Types with TypeObject
- Annex D: DDS Built-in Topic Data Types

2.2.2 Basic Network Interoperability Profile

This profile adds type safety to the Minimal profile. It enables checking type compatibility between published and subscribed types as a precondition for matching the endpoints.

Conformance with the Basic Network Interoperability Profile requires conformance with the Minimal Network Interoperability profile and the following clauses:

- “Built-in TypeLookup service” (Clause 7.6.2.3)

2.3 Optional XTYPES 1.1 Interoperability Profile

This profile adds interoperability with implementations that conform with version 1.1 of the XTYPES specification.

Conformance with the XTYPES 1.1 Interoperability Profile requires conformance with the Basic Network Interoperability profile and support of version 1 encoding in Clause 7.6.2.1.1 “DataRepresentationQosPolicy: Conceptual Model.”

2.4 Optional XML Data Representation Profile

This profile adds support for the XML Data Representation format.

Conformance to this profile requires conformance to the following clauses of this specification:

- “XML Type Representation” (Clause 7.3.2)
- “XSD Type Representation” (Clause 7.3.3)
- “XML Data Representation” (Clause 7.4.4)
- The XML schemas defined by Annex A: XML Type Representation Schema

3. Normative References

The following normative documents contain provisions that, through reference in this text, constitute provisions of this specification.

- **[DDS]** *Data Distribution Service for Real-Time Systems Specification*, Version 1.2 (OMG document formal/2007-01-01)
- **[RTPS]** *Real-Time Publish-Subscribe Wire Protocol DDS Interoperability Wire Protocol Specification*, Version 2.22 (OMG document formal/2014-09-01)
- **[DDS-XTYPES11]** *Extensible and Dynamic Topic Types for DDS Specification*, Version .1.1 (OMG document formal/2014-11-03)
- **[IDL41]** *Interface Definition Language*, Version 4.1 (OMG document ptc/16-11-11)
- **[CDR]** *Common Object Request Broker Architecture (CORBA) Specification*, Version 3.1, Part 2 (OMG document formal/2008-01-07), clause 9.3: “CDR Transfer Syntax”
- **[C-LANG]** *Programming languages -- C* (ISO/IEC document 9899:1990)
- **[C++-LANG]** *Programming languages -- C++* (ISO/IEC document 14882:2003)
- **[JAVA-LANG]** *The Java Language Specification, Second Edition* (by Sun Microsystems, <http://java.sun.com/docs/books/jls/>)
- **[C-MAP]** *C Language Mapping Specification*, Version 1.0 (OMG document formal/1999-07-35)
- **[C++-MAP]** *C++ Language Mapping Specification*, Version 1.2 (OMG document formal/2008-01-09)
- **[JAVA-MAP]** *IDL to Java Language Mapping*, Version 1.3 (OMG document formal/2008-01-11)
- **[DDS-PSM-CXX]** *ISO/IEC C++ 2003 Language DDS PSM™*, Version 1.0 (OMG document formal/2013-11-01)
- **[IDL-XSD]** *CORBA to WSDL/SOAP Interworking Specification*, Version 1.2.1 (OMG document formal/2008-08-03)
- **[LATIN]** *Information technology -- 8-bit single-byte coded graphic character sets -- Part 1: Latin alphabet No. 1* (ISO/IEC document 8859-1:1998)
- **[UCS]** *Information technology -- Universal Multiple-Octet Coded Character Set (UCS)* (ISO/IEC document 10646:2003)
- **[FNMATCH]** *POSIX fnmatch function* (IEEE 1003.2-1992 clause B.6)
- **[ISO-8601:2004]** *ISO 8601:2004 1988 (E), "Data elements and interchange formats - Information interchange - Representation of dates and times"*.
- **[IETF RFC 3339]** *IETF RFC 3339, "Date and Time on the Internet: Timestamps"*.
<https://tools.ietf.org/html/rfc3339>.

- **[UNICODE]** *The Unicode Standard, Version 9.0.0.* (Mountain View, CA: The Unicode Consortium, 2016. ISBN 978-1-936213-13-9). <http://www.unicode.org/versions/Unicode9.0.0/>.
- **[IEEE-754]** IEEE Standard for Binary Floating-Point Arithmetic, 754-2008 - IEEE Standard for Floating-Point Arithmetic

4. Terms and Definitions

Data Centric Publish-Subscribe (DCPS) – The mandatory portion of the DDS specification used to provide the functionality required for an application to publish and subscribe to the values of data objects.

Data Distribution Service (DDS) – An OMG distributed data communications specification that allows Quality of Service policies to be specified for data timeliness and reliability. It is independent of implementation languages.

5. Symbols

No additional symbols are used in this specification.

6. Additional Information

6.1 Data Distribution Service for Real-Time Systems (DDS)

The *Data Distribution Service for Real-Time Systems* (DDS) is the Object Management Group (OMG) standard for data-centric publish-subscribe communication. This standard has experienced a record-pace adoption within the Aerospace and Defense domain and is swiftly expanding to new domains, such as Transportation, Financial Services, and SCADA. To sustain and further propel its adoption, it is essential to extend the DDS standard to effectively support a broad set of use cases.

The OMG DDS specification has been designed to effectively support statically defined data models. This assumption requires that the data types used by DDS Topics are known at compile time and that every member of the DDS global data space *agrees* precisely on the same topic-type association. This model allows for good properties such as static type checking and very efficient, low-overhead, implementation of the standard. However it also suffers a few drawbacks:

- It is hard to cope with data models evolving over time unless all the elements of the system affected by that change are upgraded consistently. For example, the addition or removal of a field in the data type would not possible unless all the components in the system that use that data type are upgraded simultaneously.
- Applications using a data type must know the details of the data type at compile time, preventing use cases that would require dynamic discovery of the data types and their manipulation without compile-time knowledge. For example, a data-visualization tool cannot discover dynamically the type of a particular topic and extract the data for presentation in an interface.

With the increasing adoption of DDS for the integration of large distributed systems, it is desirable to provide a mechanism that supports evolving the data types without requiring all components using that type to be upgraded simultaneously. Moreover it is also desirable to provide a “dynamic” API that allows type definition, as well as publication and subscription data types without compile-time knowledge of the schema.

Most of the concerns outlined in Scope above (Type System, Type Representation, etc.) are already addressed in the DDS specification and/or in the DDS Interoperability Protocol specification. However, these specifications sometimes are not sufficiently explicit, complete, or flexible with regards to the above concerns of large dynamic systems. This specification addresses those limitations.

The current mechanisms used by the existing specifications are shown in Table 1 below.

Table 1 – Type-related concerns addressed by this specification

<i>Concern</i>	<i>Mechanism currently in use by DDS and the Interoperability Protocol</i>
Type System	The set of “basic” IDL types: primitive types, structures, unions, sequences, and arrays. This set is only implicitly defined.
Type Representation	Uses OMG Interface Definition language (IDL). This format is used to describe types on a file. There is no representation provided for communication of types over the network.
Data Representation	The DDS Interoperability Protocol uses the OMG Common Data Representation (CDR) based on the corresponding IDL type. It also uses a “parameterized” CDR representation for the built-in Topics, which supports schema evolution.
Language Binding	Plain Language objects as defined by the IDL language mapping.

This specification formally addresses each of the aforementioned concerns and specifies multiple mechanisms to address each concern. Multiple mechanisms are required to accommodate a broad range of application requirements and balance tradeoffs such as efficiency, evolvability, ease of integration with other technologies (such as Web Services), as well as compatibility with deployed systems. Care has been taken such that the introduction of multiple mechanisms does not break existing systems nor make it harder to develop future interoperable systems.

Table 2 summarizes the main features and mechanisms provided by the specification to address each of the above concerns.

Table 2 – Main features and mechanisms provided to address type-related concerns

<i>Concern</i>	<i>Features and mechanisms introduced by the Extensible Topics specification</i>
Type System	Defined in UML, independent of any programming language. Supports: <ul style="list-style-type: none"> • Most of the IDL data types • Specification of additional DDS-specific concepts, such as keys • Single Inheritance • Type versioning and evolution • Sparse types (types, the samples of which may omit values for certain fields; see below for a formal treatment)

Type Representation	<p>Several specified:</p> <ul style="list-style-type: none"> • IDL – Supports existing IDL-defined types. • XSD – Allows reuse of schemas defined for other purposes (e.g., in WSDL files). • XML – Provides a compact, XML-based representation suitable for human input and tool use. • TypeObject – The most compact representation (typically binary). Optimized for network propagation of types.
Data Representation	<p>Several specified:</p> <ul style="list-style-type: none"> • CDR – Most compact representation. Binary. Interoperates with existing systems. Does not support evolution. • Parameterized CDR – Binary representation that supports evolution. It is the most compact representation that can support type evolution. • XML – Human-readable representation that supports evolution.
Language Binding	<p>Several Specified:</p> <ul style="list-style-type: none"> • Plain Language Binding – Equivalent to the type definitions generated by existing standard IDL language mappings. Convenient. Requires compile-time knowledge of the type. • Dynamic Language Binding – Allows dynamic type definition and introspection. Allows manipulation of data without compile-time knowledge.

6.2 Acknowledgments

The following companies submitted and/or supported parts of this specification:

- Real-Time Innovations
- PrismTech Corp
- THALES
- Twin Oaks Computing, Inc.
- Object Computing, Inc.

7. Extensible and Dynamic Topic Types for DDS

7.1 Overview

A running DDS [DDS] application that publishes and subscribes data must deal directly or indirectly with data types and data samples of those types and the various representations of those objects. The application and middleware perspectives related to data and data types are shown in Figure 2 below.

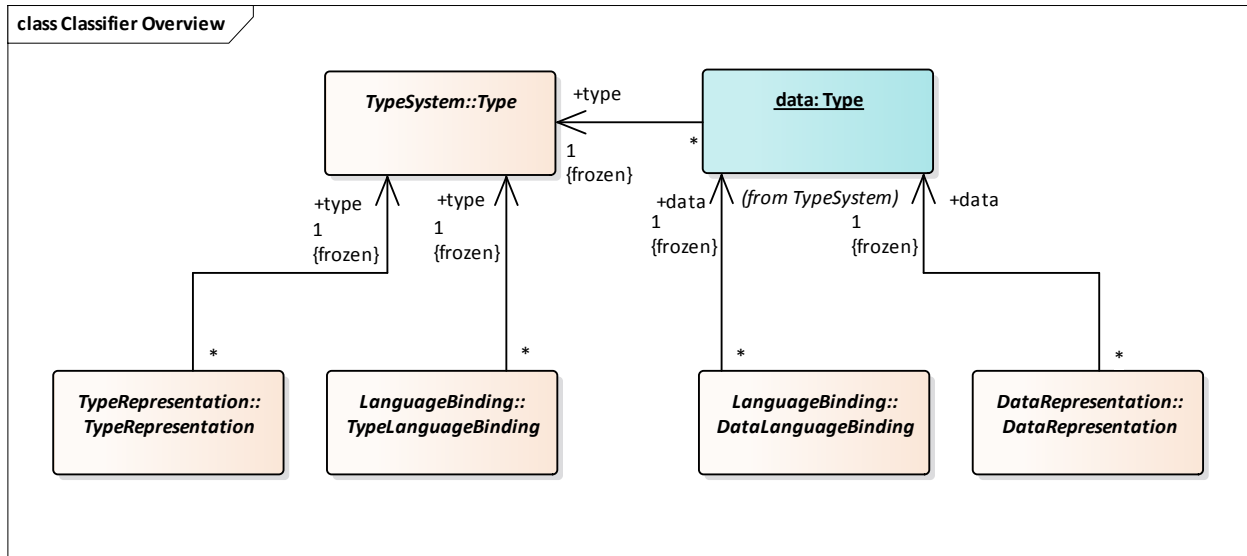


Figure 2 – Relationships between Type System, Type Representation, Language Binding, and Data Representation

DDS data objects have an associated data type (in the common programming language sense of the word) that defines a common structure for all objects of the type. From a programming perspective, an object is manipulated using a Language Binding suitable for the programming language in use (e.g., Java). From a network communications and file storage perspective, an object must have a representation (encoding) that is platform neutral and maps into a contiguous set of bytes, whether textual or binary.

Similarly, from a programming perspective a data type is manipulated using a Language Binding to the programming language of choice (sometimes known as a reflection API) and must have a representation (encoding) that is platform neutral and maps into a contiguous set of bytes (e.g., XSD or IDL).

The following example is based on a hypothetical “Alarm” data use case can be used to explain Figure 2 above.

An application concerned with alarms might use a type called “AlarmType” to indicate the nature of the alarm, point of origin, time when it occurred, severity, etc. Applications publishing and subscribing to AlarmType must therefore understand to some extent the logical or semantic contents of that type. This is what is represented by the `TypeSystem::Type` class in Figure 2 above.

If this type is to be communicated in a design document or electronically to a tool, it must be represented in some “external” format suitable for storing in a file or on a network packet. This aspect is represented by the `TypeRepresentation::TypeRepresentation` class in Figure 2 above. A realization of the `TypeRepresentation` class may use XML, XSD, or IDL to represent the type.

An application wishing to understand the structure of the `Type`, or the middleware attempting to check type-compatibility between writers and readers, must use some programming language construct to examine the type. This is represented by the `LanguageBinding::TypeLanguageBinding` class. As an example of this concept, the class `java.lang.Class` plays this role within the Java platform.

An application publishing Alarms or receiving Alarms must use some programming language construct to set the value of the alarm or access those values when it receives the data. This programming language construct may be a plain language object (such as the one generated from an IDL description of the type) or a dynamic container that allows setting and getting named fields, or some other programming language object. This is represented by the `LanguageBinding::DataLanguageBinding` class.

An application wishing to store Alarms on a file or the middleware wishing to send Alarms on a network packet or create Alarm objects from data received on the network must use some mechanism to “serialize” the Alarm into bytes in a platform-neutral fashion. This is represented by the `DataRepresentation::DataRepresentation` class. An example of this would be to use the CDR Representation derived from the IDL Type Representation.

The classes in Figure 2 above represent each of the independent concerns that both application and middleware need to address. The non-normative Figure 3 below indicates their relationships to one another in a less formal way.

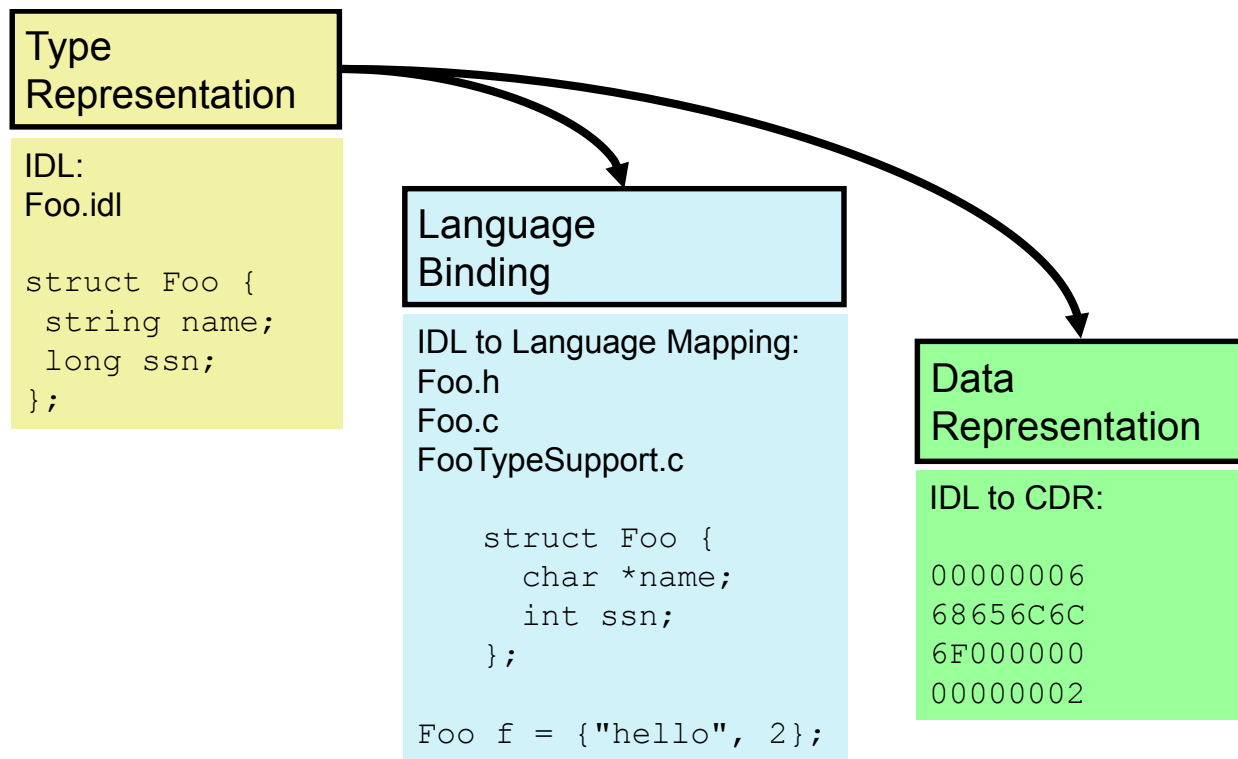


Figure 3 – Example Type Representation, Language Binding, and Data Representation

Type Representation is concerned with expressing the type in a manner suitable for human input and output, file storage, or network communications. IDL is an example of a standard type representation. Language Binding is concerned with the programming language constructs used to interact with data of a type or to introspect the type. Plain language objects as obtained from the IDL language mappings of the IDL representation of the type are one possible Language Binding. Data Representation is concerned with expressing the data in a way that can be stored in a file or communicated over a network or manipulated by a human. The Common Data Representation is a Data Representation optimized for network communications; XML is another representation more suitable for human manipulation.

7.2 Type System

The Type System defines the data types that can be used for DDS Topics and therefore the type of the data that can be published and subscribed via DDS.

7.2.1 Background (Non-Normative)

The specified type system is designed to be sufficiently rich to encompass the needs of modern distributed applications and cover the basic data types available both in common programming languages such as C/C++, Java, and C#, as well as in distributed computing data-definition languages such as IDL or XDR.

The specified type system supports the following primitive types:

- Boolean type
- Byte type
- Integral types of various bit lengths (16, 32, 64), both signed and unsigned
- Floating point types of various precisions: single precision, double precision, and quad precision
- Single-byte and wide character types

In addition the specified type system covers the following non-basic types constructed as collections or aggregations of other types:

- Structures, which can singly inherit from other structures
- Unions
- Single- and multi-dimensional arrays
- Variable-length sequences of a parameterized element type
- Strings of single-byte and wide characters
- Variable-length maps of parameterized key and value types

The specified type-system supports type evolution, type inheritance, and sparse types. These concepts are described informally in Clauses 7.2.1.1, 7.2.1.2, and 7.2.1.3 below and formally in Clause 7.2.2.

7.2.1.1 Type Evolution Example

Assume a DDS-based distributed application has been developed that uses the Topic “Vehicle Location” of type `VehicleLocationType`. The type `VehiclePositionType` itself was defined using the following IDL:

```
// Initial Version
struct VehicleLocationType {
    float latitude;
    float longitude;
};
```

As the system evolves it is deemed useful to add additional information to the `VehicleLocationType` such as the estimated error latitude and longitude errors as well as the direction and speed resulting in:

```
// New version
struct VehicleLocationType {
    float latitude;
    float longitude;
    float latitude_error_estimate;    // added field
    float longitude_error_estimate;  // added field
```

```

float direction;           // added field
float speed;              // added field
};

```

This additional information can be used by the components that understand it to implement more elaborate algorithms that estimate the position of the vehicle between updates. However, not all components that publish or subscribe data of this type will be upgraded to this new definition of `VehicleLocationType` (or if they will not be upgraded, they will not be upgraded at the same time) so the system needs to function even if different components use different versions of `VehicleLocationType`.

The Type System supports type evolution so that it is possible to “evolve the type” as described above and retain interoperability between components that use different versions of the type such that:

- A publisher writing the “initial version” of `VehicleLocationType` will be able to communicate with a subscriber expecting the “new version” of the `VehicleLocationType`. In practice what this means is that the subscriber expecting the “new version” of the `VehicleLocationType` will, depending on the details of how the type was defined, either be supplied some default values for the added fields or else be told that those fields were not present.
- A publisher writing the “new version” of `VehicleLocationType` will be able to communicate with a subscriber reading the “initial version” of the `VehicleLocationType`. In practice this means the subscriber expecting the “initial version” of the `VehicleLocationType` will receive data that strips out the added fields.

Evolving a type requires that the designer of the new type explicitly tags the new type as equivalent to, or an extension of, the original type and limits the modifications of the type to the supported set. The addition of new fields is one way in which a type can be evolved. The complete list of allowed transformations is described in Clause 7.2.4.

7.2.1.2 Type Inheritance Example

Building upon the same example in Clause 7.2.1.1, assume that the system that was originally intended to only monitor location of land/sea-surface vehicles is now extended to also monitor air vehicles. The location of an air vehicle requires knowing the altitude as well. Therefore the type is extended with this field.

```

// Extended Location
struct VehicleLocation3DType : VehicleLocationType {
    float altitude;
    float vertical_speed;
};

```

`VehicleLocation3DType` is an extension of `VehicleLocationType`, not an evolution. `VehicleLocation3DType` represents a new type that extends `VehicleLocationType` in the object-oriented programming sense (IS-A relationship).

The Type System supports type inheritance so that it is possible to “extend the type” as described above and retain interoperability between components that use different versions of the type. So that:

- An application subscribing to Topic “Vehicle Position” and expecting to read `VehicleLocationType` CAN receive data from a Publisher that is writing a `VehicleLocation3DType`. In other words applications can write extended types and read base types.
- An application subscribing to Topic “Vehicle Position” and expecting to read `VehicleLocation3DType` CAN receive data from a Publisher that is writing a `VehicleLocationType`. Applications expecting the derived (extended) type can accept the base type; additional members in the derived type will take no value or a default value, depending on their definitions.

This behavior matches the behavior of the “IS-A” relationship in Object-Oriented Languages,

Intuitively this means that a `VehicleLocation3DType` is a new type that happens to extend the previous type. It can be substituted in places that expect a `VehiclePosition` but it is not fully equivalent. The substitution only works one way: An application expecting a `VehicleLocation3DType` cannot accept a `VehiclePosition` in place because it cannot “just” assume some default value for the additional fields. Rather it wants to just read those `VehiclePosition` that corresponds to Air vehicles.

7.2.1.3 Sparse Types Example

Suppose that an application publishes a stream of events. There are many kinds of events that could occur in the system, but they share a good deal of data, they must all be propagated with the same QoS, and the relative order among them must be preserved—it is therefore desirable to publish all kinds of events on a single topic. However, there are fields that only make sense for certain kinds of event. In its local programming language (say, C++ or Java), the application can assign a pointer to null to omit a value for these fields. It is desirable to extend this concept to the network and allow the application to omit irrelevant data in order to preserve the correct semantics of the data.

Alternatively, suppose that an application subscribes to data of a type containing many fields, most of which often take a pre-specified “default value” but may, on occasion, deviate from that default. In this situation it would be inefficient to send every field along with every sample. Rather it would be better to just send the fields that take a non-default value and fill the missing fields on the receiving side, or even let the receiving application do that job. This situation occurs, for example, in the DDS Built-in Topic Data. It also occurs in financial applications that use the FIX encoding for the data.

The type system supports sparse types whereby a type can have fields marked “optional” so that a Data Representation may omit those fields. Values for non-optional fields may also be omitted to save network bandwidth, in which case the Service will automatically fill in default values on behalf of the application.

7.2.2 Type System Model

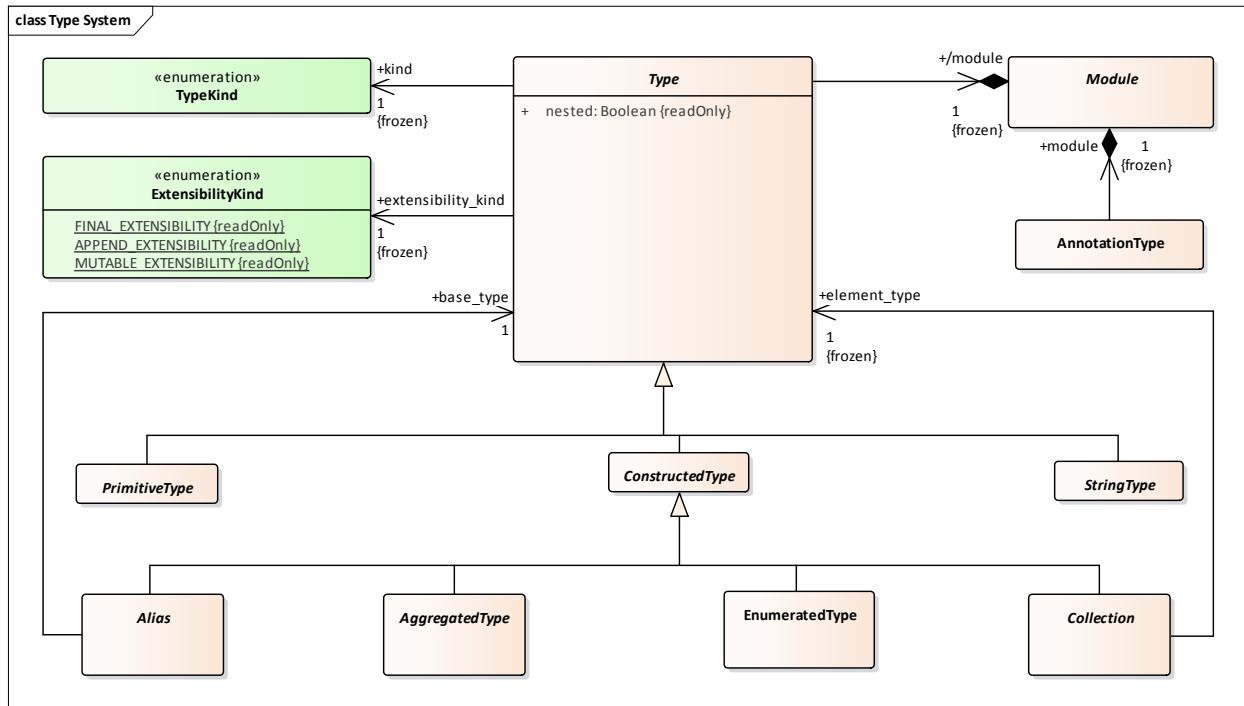


Figure 4 – Type System Model

The definition of a type in the Type System can either be primitive or it can be constructed from the definitions of other types.

The Type System model is shown in Figure 4. This model has the following characteristics:

- A type has a non-empty *name* that is unique within its namespace (see Clause 7.2.2.1). The set of valid names is the set of valid identifiers defined by the OMG IDL specification [IDL41].
- A type has a *kind* that identifies which primitive type it is or, if it is a constructed type, whether it is a structure, union, sequence, etc.
- The type system supports Primitive Types (i.e., their definitions do not depend on those of any other types) whose names are predefined. The Primitive Types are described in 7.2.2.2.
- The type system supports Constructed Types whose names are explicitly provided as part of the type-definition process. Constructed Types include enumerations, collections, structure, etc. Constructed types are described in Clause 7.2.2.4.

7.2.2.1 Namespaces

A namespace defines the scope within which a given name must be unique. That is, it is an error for different elements within the same namespace to have the same name. However, it is legal for different elements within different namespaces to have the same name.

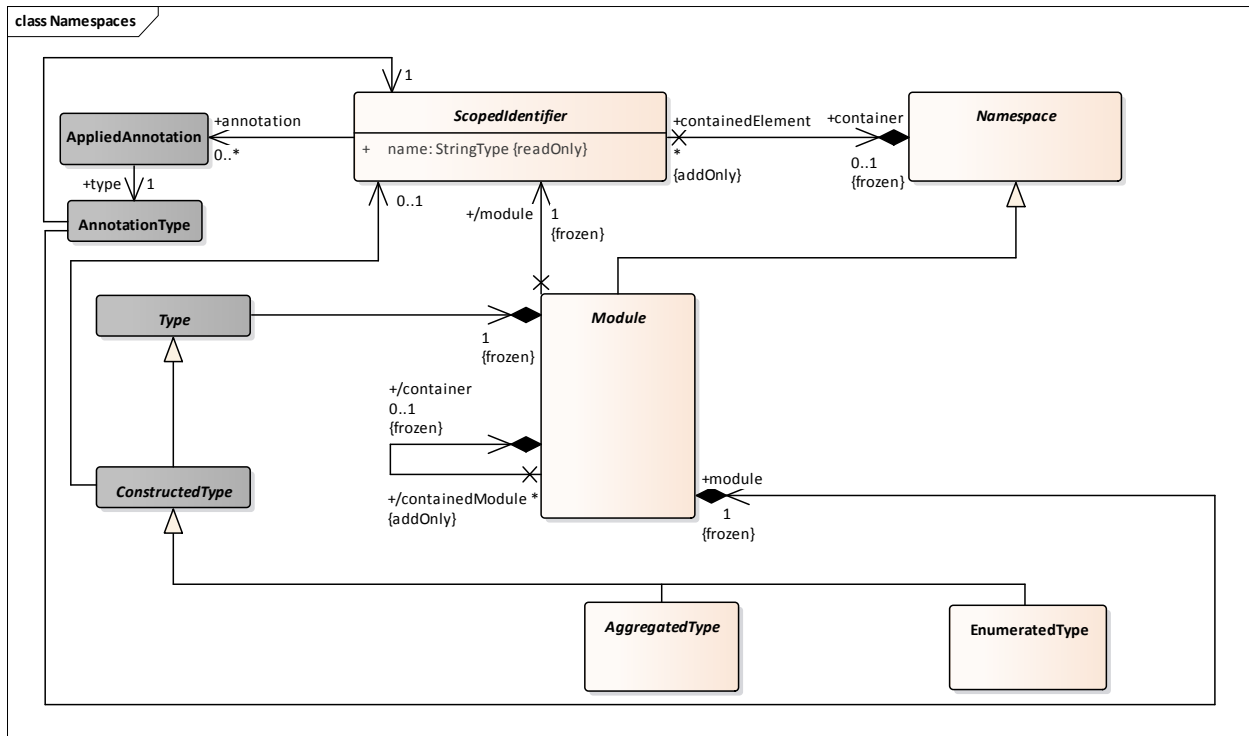


Figure 5 – Namespaces

Namespaces fall into one of two categories:

- *Modules* are namespaces whose contained named elements are types. The concatenation of module names with the name of a type inside of those modules is referred to as the type’s “fully qualified name.”
- Certain kinds of *types* are themselves namespaces with respect to the elements inside of them.

7.2.2.2 Primitive Types

The primitive types in the Type System have parallels in most computer programming languages and are the building blocks for more complex types built recursively as collections or aggregations of more basic types.

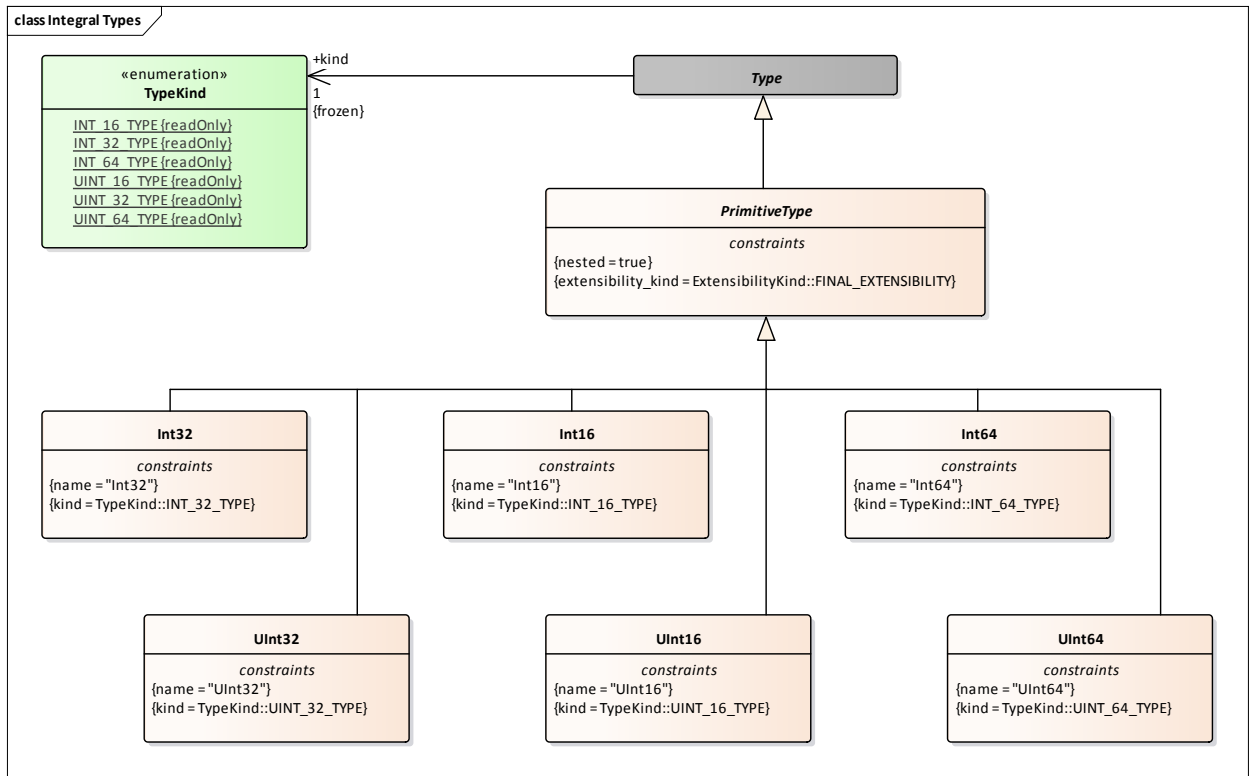


Figure 6 – Primitive Types: Integral Types

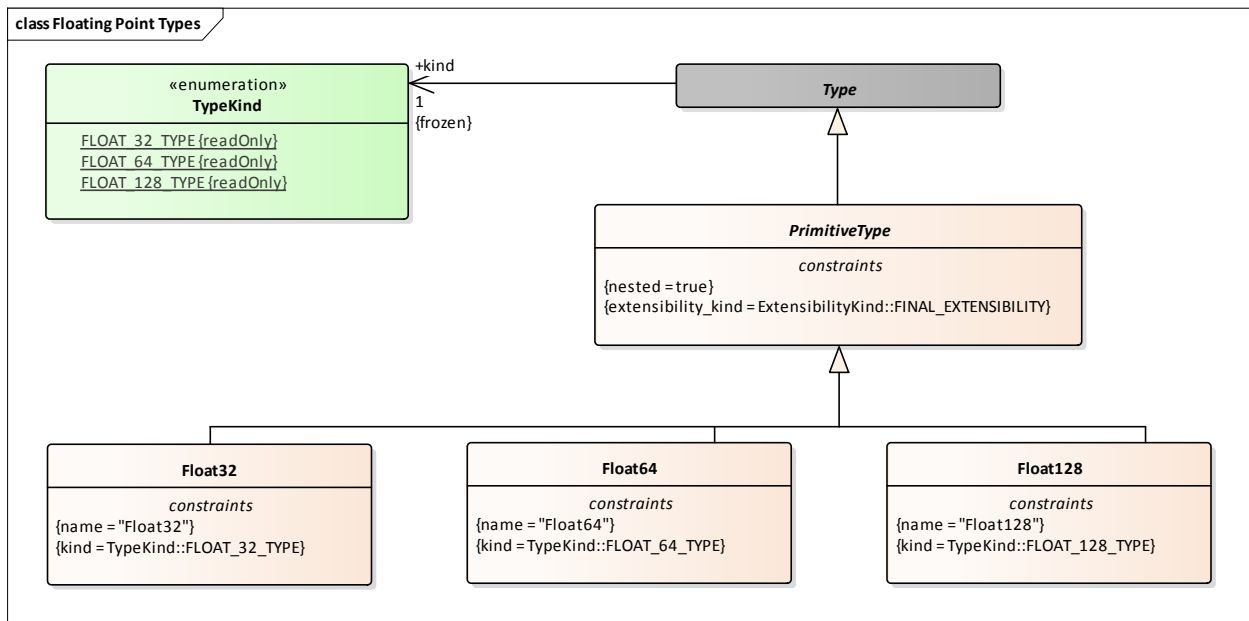


Figure 7 – Primitive Types: Floating Point Types

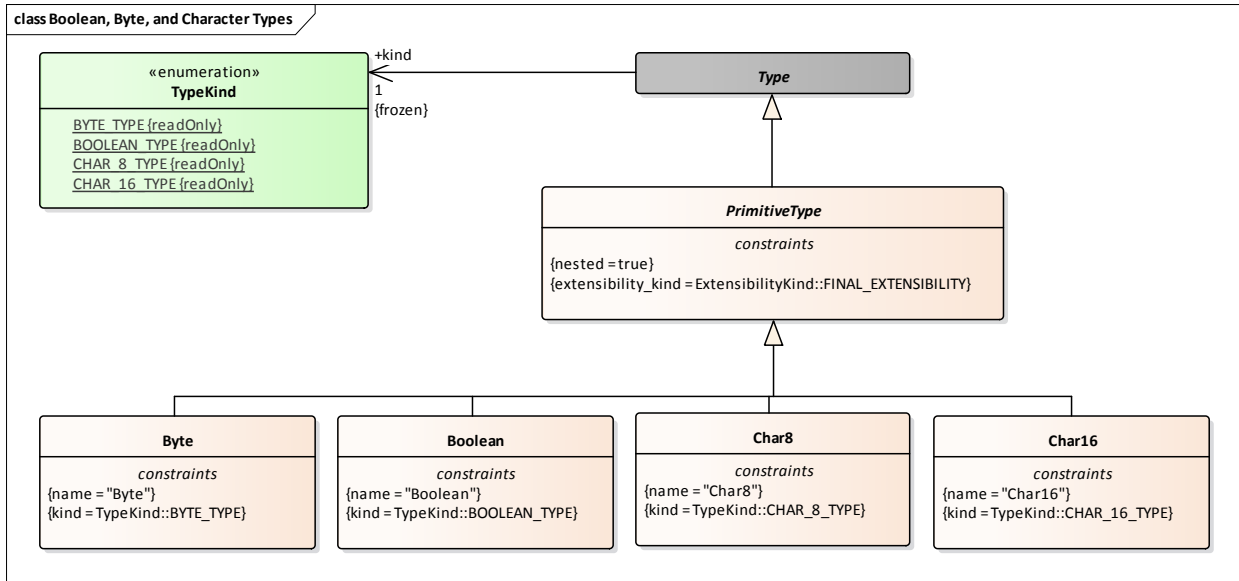


Figure 8 – Primitive Types: Booleans, Bytes, and Characters

Primitive Types include the primitive types present in most programming languages, including Boolean, integer, floating point, and character.

Table 3 below enumerates and describes the available primitive types. Note that value ranges are in this package specified only in terms of upper and lower bounds; data sizes and encodings are the domain of the Type Representation and Data Representation packages.

Table 3 – Primitive Types

<i>Type Kind</i>	<i>Type Name</i>	<i>Description</i>
BOOLEAN_TYPE	Boolean	Boolean type. Data of this type can only take two values: true and false.
BYTE_TYPE	Byte	Single opaque byte. A Byte value has no numeric value.
INT_16_TYPE	Int16	Signed integer minimally capable of representing values in the range -32738 to +32737.
UINT_16_TYPE	UInt16	Unsigned integer minimally capable of representing values in the range 0 to +65535.
INT_32_TYPE	Int32	Signed integer minimally capable of representing values in the range -2147483648 to +2147483647.
UINT_32_TYPE	UInt32	Unsigned integer minimally capable of representing values in the range 0 to +4294967295.
INT_64_TYPE	Int64	Signed integer minimally capable of supporting values in the range -9223372036854775808 to +9223372036854775807.

UINT_64_TYPE	UInt64	Unsigned integer minimally capable of supporting values in the range 0 to +18446744073709551617.
FLOAT_32_TYPE	Float32	Floating point number minimally capable of supporting the range and precision of an IEEE 754 single-precision floating point value.
FLOAT_64_TYPE	Float64	Floating point number minimally capable of supporting the range and precision of an IEEE 754 double-precision floating point value.
FLOAT_128_TYPE	Float128	Floating point number minimally capable of supporting the range and precision of an IEEE 754 quadruple-precision floating point value.
CHAR_8_TYPE	Char8	8-bit character type. There is no encoding specified, it may be ASCII, ISO-8859-1, or used to hold a byte of a multi-byte-encoded character set.
CHAR_16_TYPE	Char16	16-bit character type capable of supporting the Basic Multilingual Plane (BMP) encoded in UTF-16.

The primitive types do not exist within any module; their names are top-level names.

7.2.2.2.1 Character Data

The character types identified above require further definition, provided here.

7.2.2.2.1.1 Design Rationale (Non-Normative)

Because the Unicode character set is a superset of the US-ASCII character set, some readers may question why this specification provides two types for character data: `Char8` and `Char16`. These types are differentiated to facilitate the efficient representation and navigation of character data as well as to more accurately describe the designs of existing systems.

Existing languages for type definition—including C, C++, and IDL—distinguish between regular and wide characters (C/C++ `char` vs. `wchar_t`; IDL `char` vs. `wchar`). While other commonly used typing systems do not make such a distinction—in particular Java and the ECMA Common Type System, of which Microsoft’s .Net is an implementation—it is more straightforward to map two platform-independent types to a single platform-specific type than it is to map objects of a single platform-independent type into different platform-specific types based on their values.

7.2.2.2.1.2 Character Sets and Encoding

7.2.2.2.1.2.1 Use of Unicode

This specification uses the Unicode Standard (version 9.0, June 2016) as the means to represent characters and strings.

Unicode defines a codespace of 1,114,112 code points in the range 0x000000 to 0x10FFFF. A Unicode code point is referred to by writing "U+" followed by its hexadecimal number (e.g. U+0000F1).

In the Unicode standard, a plane is a continuous group of 2^{16} code points. There are 17 planes, identified by the numbers 0 to 16, which corresponds with the possible values 0x00-0x10 of the first two positions in six position format (hhhhhh).

Plane 0 is called the Basic Multilingual Plane (BMP). It contains nearly all commonly used writing systems and symbols. It contains characters U+0000 to U+FFFF. Planes 1–16, are called “supplementary planes”. As of Unicode version 9.0, six of the planes have assigned code points (characters), and four are named.

Unicode can be implemented by different character encodings. The most commonly used encodings are UTF-8, UTF-16, and UTF-32 (in that order). The Unicode code point is shared across all these encodings.

The UTF-8 encoding is backward compatible with the ASCII character set and is the default one used by most C and C++ compilers. The UTF-8 representation of ASCII characters uses one 8-bit code unit. The UTF-8 representation ISO-8859-1 characters that are not in the ASCII subset uses two 8-bit code units. Any character in the Basic Multilingual Plane is encoded using one to three UTF-8 code units.

The UTF-16 encoding represents the code points in the Basic Multilingual Plane using one 16-bit code unit. The remaining Unicode characters use two 16-bit code units. The representation is numerically equal to the corresponding code points using the selected endianness.

7.2.2.2.1.2.2 CHAR_8_TYPE

This specification does not define an encoding for the `CHAR_8_TYPE`. The only constraint is that it shall be representable using 8 bits.

Rationale

By not specifying an encoding for `CHAR_8_TYPE` it is possible to use the 8-bit code-unit to either store a single ISO-8859-1 character or alternatively a code-unit of a UTF-8 encoded string.

7.2.2.2.1.2.3 Array or Sequence of CHAR_8_TYPE

This specification does not define an encoding for the `CHAR_8_TYPE` that appears as an element of an array or sequence of `CHAR_8_TYPE`.

Rationale

By not specifying an encoding for the elements of an Array or Sequence of `CHAR_8_TYPE` it becomes possible to store the characters of a String type into an Array or Sequence of `CHAR_8_TYPE` regardless of the encoding used in the String.

7.2.2.2.1.2.4 String<Char8> type

The default encoding for `String<Char8>` shall be UTF-8. This encoded shall be used for the externalized Data Representation (see clause 7.4). Language bindings (see Clause 7.5) may use

the representation that is most natural in that particular language. If this is different from UTF-8 the language binding shall manage the transformation to/from UTF-8 external Data Representation.

7.2.2.2.1.2.5 CHAR_16_TYPE

The `CHAR_16_TYPE` shall be restricted to representing Unicode codepoints in the Basic Multilingual Plane. That is Unicode codepoints from 0x0000 to U+FFFF.

The `CHAR_16_TYPE` encoding shall be UTF-16.

Rationale

UTF-16 is more space efficient than UTF-32. UTF-16 also maps directly to the Java and C# languages, which makes serialization and deserialization simple in those languages.

The BMP captures nearly all commonly used writing systems and symbols. Restricting to the BMP ensures that each codepoint is represented using a single UTF-16 code unit (16 bits)

7.2.2.2.1.2.6 Array or Sequence of CHAR_16_TYPE

The representation of each `CHAR_16_TYPE` element of an array or sequence of `CHAR_16_TYPE` shall be UTF-16 and shall be restricted to being in the Basic Multilingual Plane (Unicode codepoints from 0x0000 to U+FFFF).

7.2.2.2.1.2.7 String<Char16> type

The encoding for `String<Char16>` shall be UTF-16. This encoded shall be used for the externalized Data Representation (see Clause 7.4). Language bindings (see Clause 7.5) may use the representation that is most natural in that particular language. If this is different from UTF-8 the language binding shall manage the transformation to/from UTF-16 external Data Representation.

7.2.2.3 String Types

StringTypes are ordered one-dimensional collections of characters. StringTypes are variable-sized; objects of a given string type can have different numbers of elements (i.e., the string object's "length"). Furthermore, the length of a given string object may change between zero and the string type's "bound" (see below) over the course of its lifetime.

A string is logically very similar to a sequence. However, the element type of a string must be either `Char8` or `Char16` (or an alias to one of these); other element types are undefined. These two collections have been distinguished in order to preserve the fidelity present in common implementation programming languages and platforms.

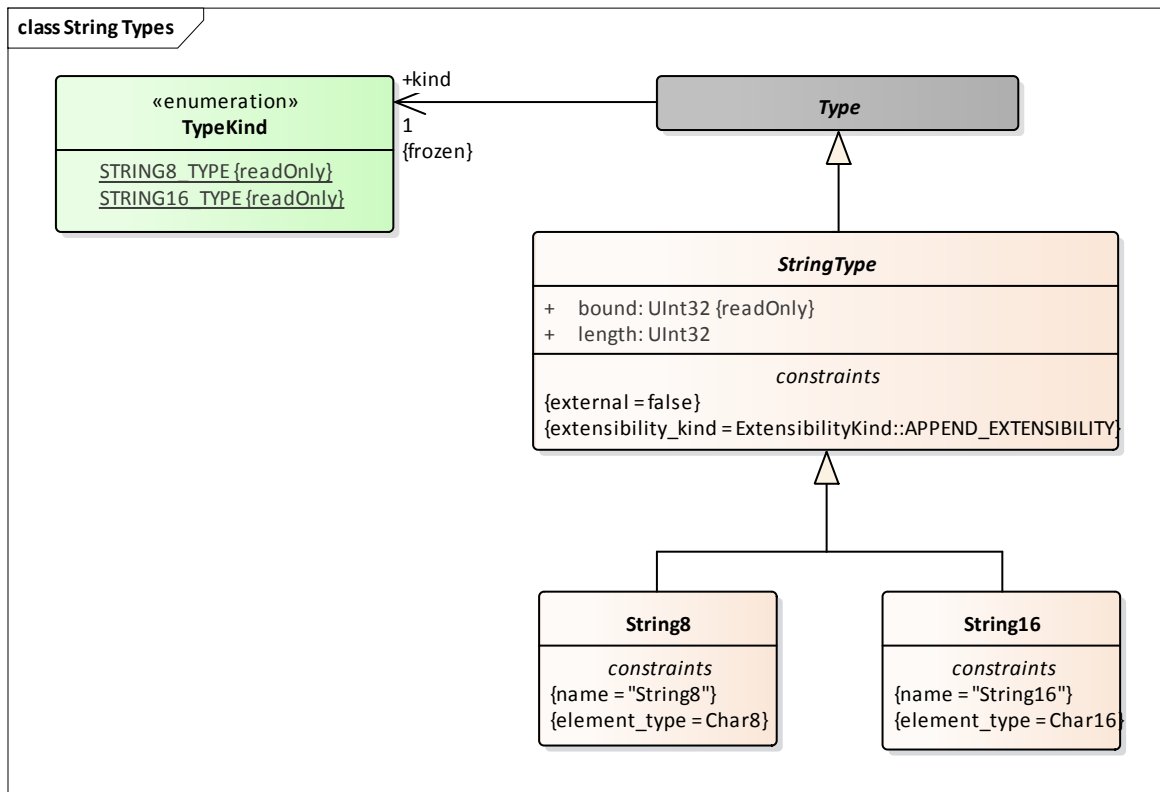


Figure 9 - String Types

7.2.2.4 Constructed Types

The definitions of these types are constructed from—that is, based upon—the definitions of other types. These other types may be either primitive types or other constructed types: type definitions may be recursive to an arbitrary depth. Constructed types are explicitly defined by a user of an implementation of this specification and are assigned a name when they are defined.

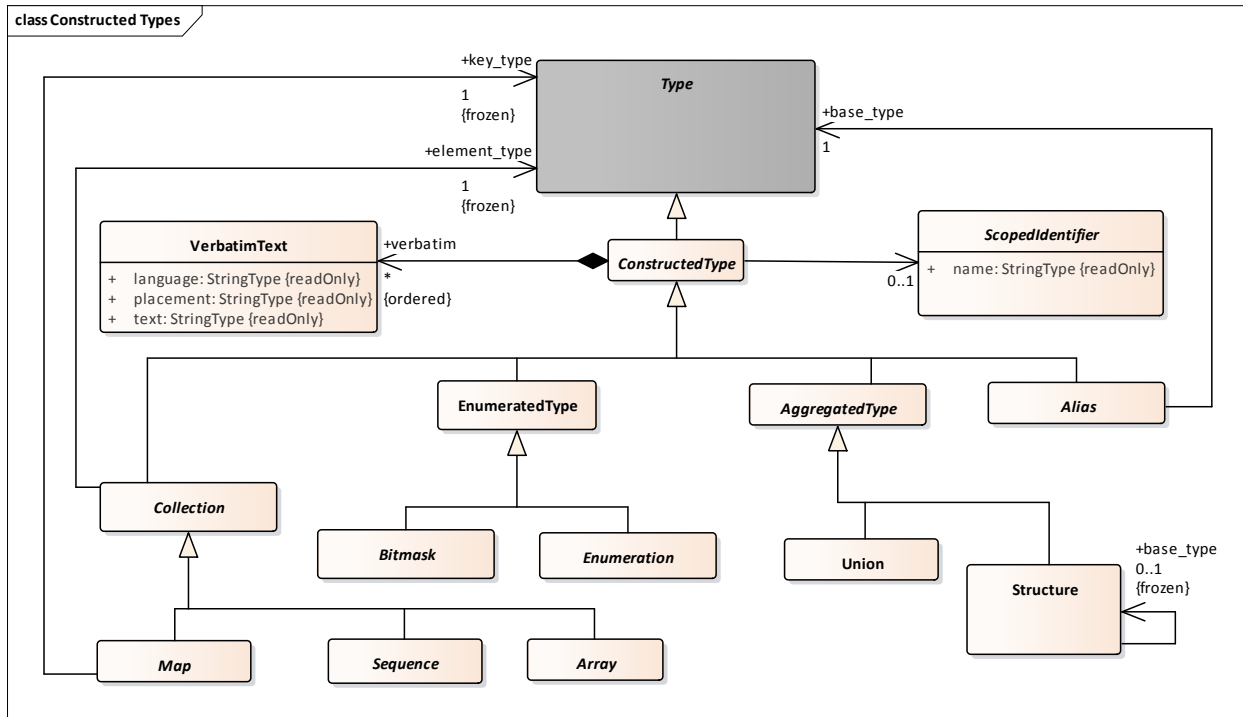


Figure 10 – Constructed Types

There are several kinds of Constructed Types: Collections, Aggregated types, Aliases, and Enumerated types. Collections are homogeneous in that all elements of the collection have the same type. Aggregated types are heterogeneous; members of the aggregated types may have different types. Aliases introduce a new name for another type. Enumerated types define a finite set of possible integer values for the data.

7.2.2.4.1 Enumerated Types

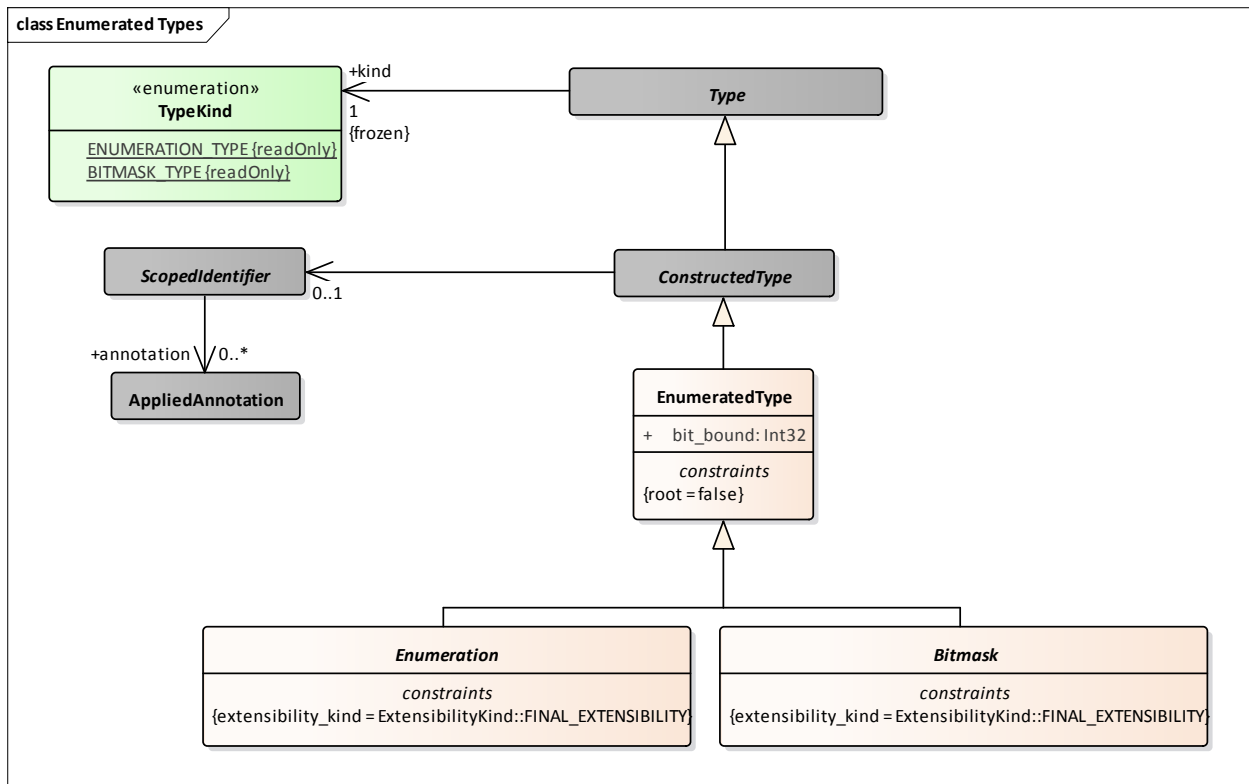


Figure 11 – Enumerated Types

7.2.2.4.1.1 Enumeration Types

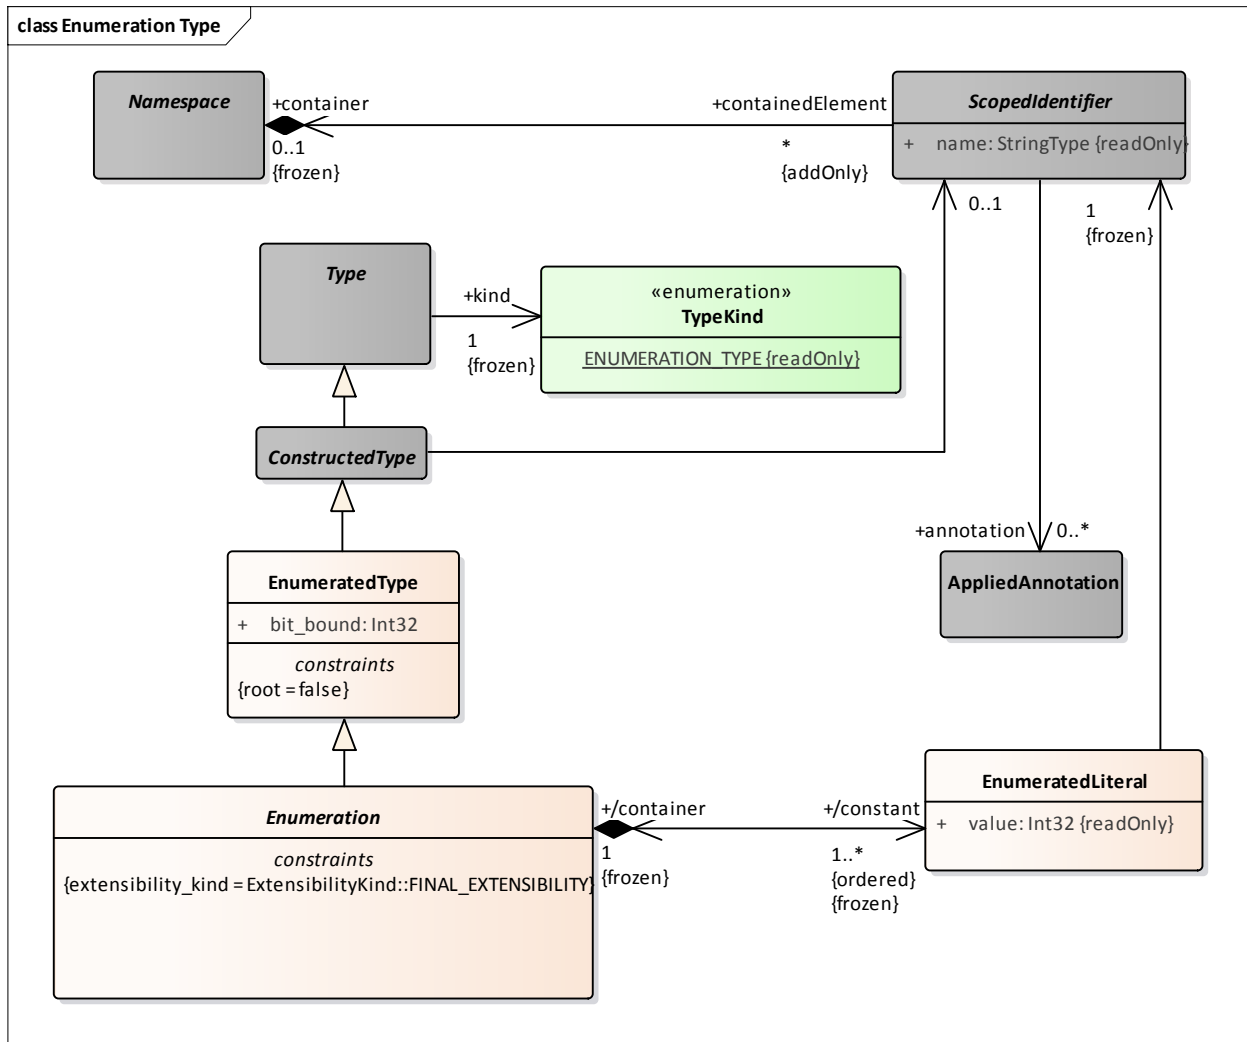


Figure 12 – Enumeration Types

Table 4 – Enumerated types

<i>Type Kind</i>	<i>Type Name</i>	<i>Description</i>
ENUMERATION_TYPE	<i>Assigned when type is defined</i>	<p>Set of literals.</p> <p>An enumerated type defines a closed set of one or more literal objects of that type. Each object of a given enumerated type has a name and an <code>Int32</code> value that are each unique within that type.</p> <p>The order in which the literals of an enumerated type are defined is significant to the definition of that type. For example, some type representations may base the numeric values of the literals on their order of definition.</p>

7.2.2.4.1.2 Bitmask Types

Bitmasks, as in the C++ standard library (and not unlike the `EnumSet` class of the Java standard library), represent a collection of Boolean flags, each of which can be inspected and/or set individually.

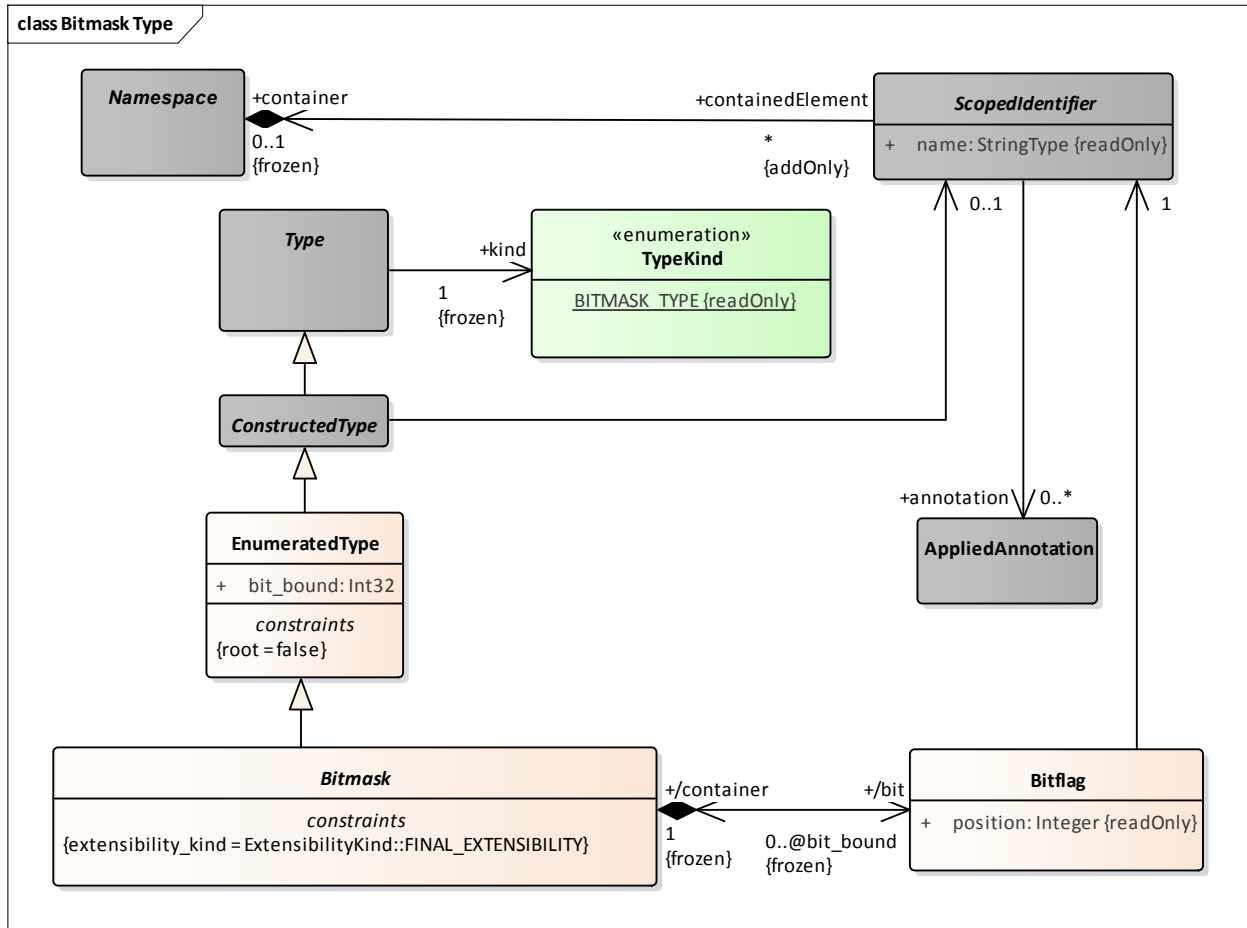


Figure 13 – Bitmask Types

Table 5 – Bitmask types

<i>Type Kind</i>	<i>Type Name</i>	<i>Description</i>
BITMASK_TYPE	<i>Assigned when type is defined</i>	<p>Ordered set of named Boolean flags.</p> <p>A bitmask defines a bound—the maximum number of bits in the set—and identifies by name certain bits within the set. The bound must be greater than zero and no greater than 64.</p>

A bitmask type reserves a number of “bits” (Boolean flags); this is referred to as its bound. (The bound of a bitmask is logically similar to the bound of an array, except that the “elements” in a bitmask are single bits.) It then identifies some subset of those bits. Each bit in this subset is

identified by name and by an index, numbered from 0 to (bound – 1). The bitmask need not identify every bit it reserves. Furthermore, the bits it does identify need not be contiguous.

Note that this type exists for the sake of semantic clarity and to enable more efficient data representations. It does not actually constrain such representations to represent each “bit” in the set as a single memory bit or to align the bitmask in any particular way.

7.2.2.4.1.2.1 Design Rationale (Non-Normative)

It is commonly the case that complex data types need to represent a number of Boolean flags. For example, in the DDS specification, status kinds are represented as `StatusKind` bits that are combined into a `StatusMask`. A bitmask (also referred to as a bit mask) allows these flags to be represented very compactly—typically as a single bit per flag. Without such a concept in the type system, type designers must choose one of two alternatives:

- Idiomatically define enumerated “kind” bits and a “mask” type. Pack and unpack the former into the latter using bitwise operators. As previously noted, this is the approach taken by the DDS specification in the case of statuses, because it predated this enhanced type model. There are several weaknesses to this approach:
 - It is verbose, both in terms of the type definition and in terms of the code that uses the bitmask; this verbosity slows understanding and can lead to programming errors.
 - It is not explicitly tied to the semantics of the data being represented. This weakness can lead to a lack of user understanding and type safety, which in turn can lead to programming errors. It furthermore hampers the development of supporting tooling, which cannot interpret the “bitmask” otherwise than as a numeric quantity.
- Represent the flags as individual Boolean values. This approach simplifies programming and provides semantic clarity. However, it is extremely verbose: a structure of Boolean members wastes at least 7/8 of the network bandwidth it uses (assuming no additional alignment and that each flag requires one bit but occupies one byte) and possible up to 31/32 of the memory it uses (on platforms such as Microsoft Windows that conventionally align Boolean values to 32-bit boundaries).

7.2.2.4.2 Alias Types

Alias types introduce an additional name for another type.

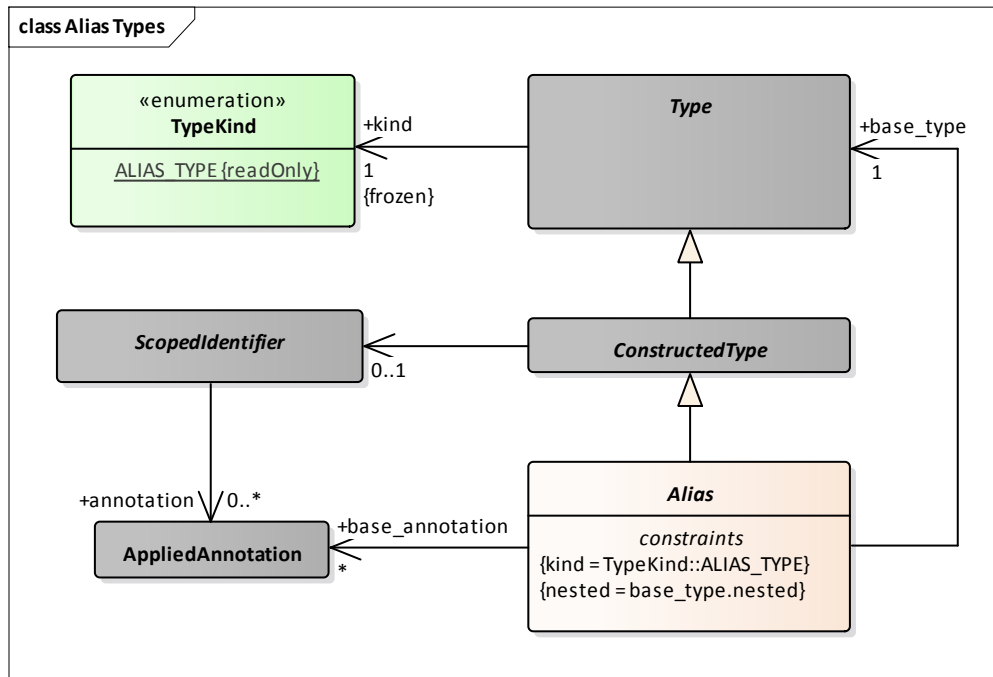


Figure 14 – Alias Types

Table 6 – Alias types

<i>Type Kind</i>	<i>Type Name</i>	<i>Description</i>
ALIAS_TYPE	<i>Assigned when type is defined</i>	<p>Alternative name for another type.</p> <p>An alias type—also referred to as a <i>typedef</i> from its representation in IDL, C, and elsewhere—applies an additional name to an already-existing type. Such an alternative name can be helpful for suggesting particular uses and semantics to human readers, making it easier to repeat complex type names for human writers, and simplifying certain language bindings.</p> <p>As in the C and C++ programming languages, an alias/typedef does not introduce a distinct type. It merely provides an alternative name by which to refer to the same type.</p>

7.2.2.4.3 Collection Types

Collections are containers for elements of a homogeneous type. The type of the element might be any other type, primitive or constructed (although some limitations apply; see below) and must be specified when the collection type is defined.

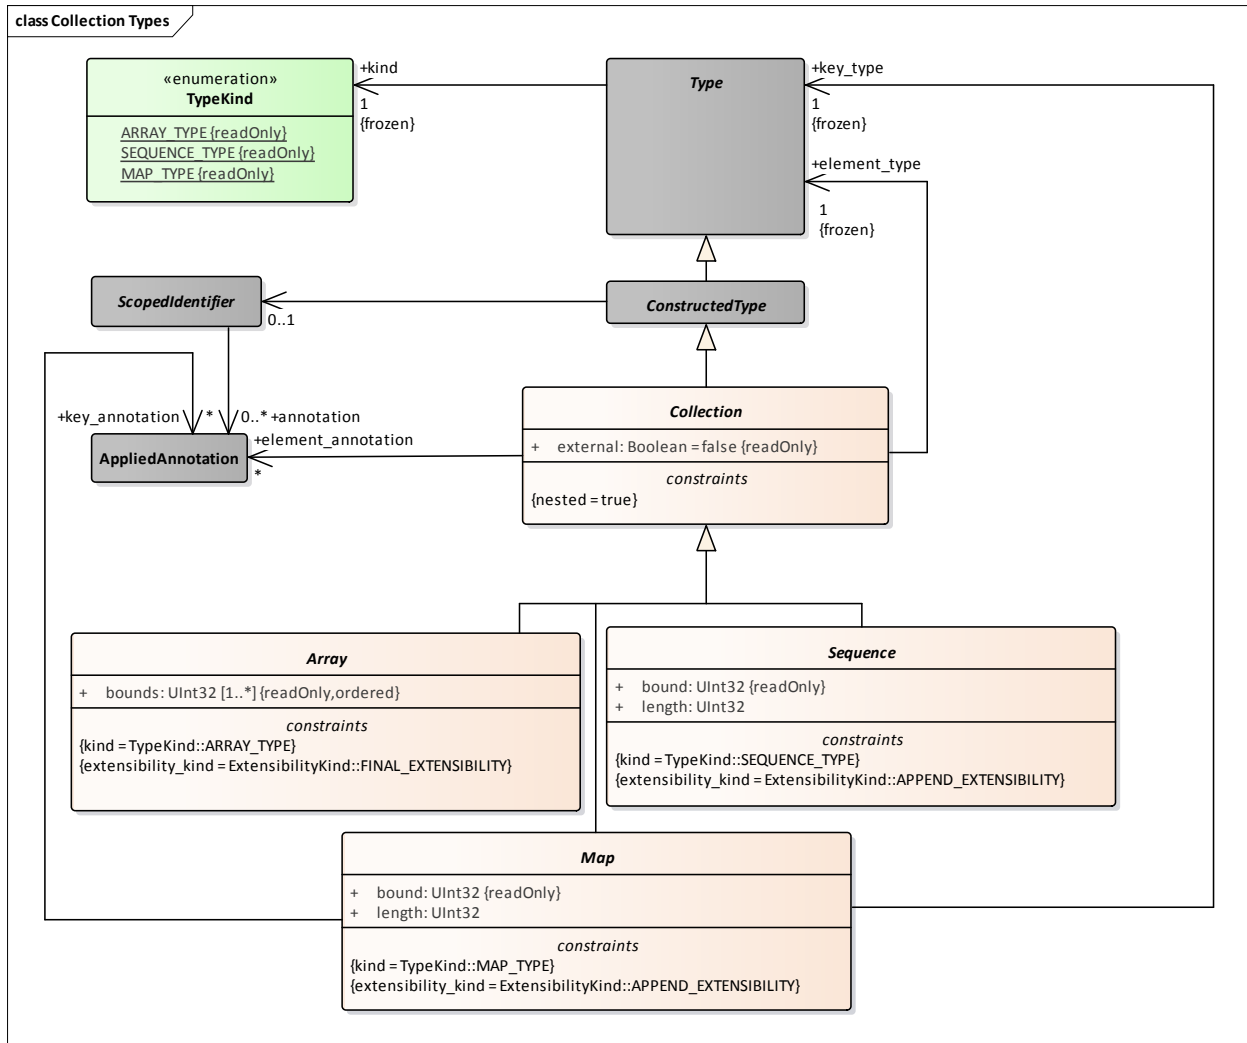


Figure 15 – Collection Types

There are three kinds of Collection Types: `ARRAY`, `SEQUENCE`, and `MAP`. These kinds are described in Table 7.

Table 7 – Collection Types

<i>Type Kind</i>	<i>Type Name</i>	<i>Description</i>
<code>ARRAY_TYPE</code>	<i>Assigned implicitly</i>	<p>Fixed-size multi-dimensional collection.</p> <p>Arrays are of a fixed size in that all objects of a given array type will have the same number of elements. Elements are addressed by a sequence of indices (one per dimension).</p> <p>Semantically, array types of higher dimensionality are distinct from arrays of arrays of lower dimensionality. (For example, a two-dimensional array is not just an array of one-dimensional arrays.) However, certain type representations may be unable to capture this distinction. (For example, IDL provides no syntax to describe an array of arrays¹, and in Java, all “multi-dimensional” arrays <i>are</i> arrays of arrays necessarily.) Such limitations in a given type representation should not be construed as a limitation on the type system itself.</p>
<code>SEQUENCE_TYPE</code>	<i>Assigned implicitly</i>	<p>Variable-size single-dimensional collection.</p> <p>Sequences are variably sized in that objects of a given sequence type can have different numbers of elements (the sequence object’s “length”); furthermore, the length of a given sequence object may change between zero and the sequence type’s “bound” (see below) over the course of its lifetime. Elements are addressed by a single index.</p>
<code>MAP_TYPE</code>	<i>Assigned implicitly</i>	<p>Variable-size associative collection.</p> <p>Maps are variably sized in that objects of a given map type can have different numbers of elements (the map object’s “length”); furthermore, the length of a given map object may change between zero and the map type’s “bound” (see below) over the course of its lifetime.</p> <p>“Map value” elements are addressed by a “map key” object, the value of which must be unique within a given map object. The types of both of these are homogeneous within a given map type and must be specified when the map type is defined.</p>

¹ An intermediate alias can help circumvent this limitation; see below for a more formal treatment of aliases.

Collection types are defined implicitly as they are used. Their definitions are based on three attributes:

- **Collection kind:** The supported kinds of collections are identified in **Table 7** above.
- **Element type:** The concrete type to which all elements conform. (Collection elements that are of a subtype of the element type rather than the element type itself may be truncated when they are serialized into a Data Representation.)

In the case of a map type, this attribute corresponds to the type of the *value* elements. Map types have an additional attribute, the *key element type*, that indicates the type of the key elements. Implementers of this specification need only support key elements of signed and unsigned integer types and of narrow and wide string types; the behavior of maps with other key element types is undefined and may not be portable. (**Design rationale, non-normative:** Support for arbitrary key element types would require implementers to provide uniform sorting and/or hashing operations, which would be impractical on many platforms. In contrast, these operations have straightforward implementations for integer and string types.)

- **Bound:** The maximum number of elements the collection may contain (inclusively); it must be greater than zero.

In the cases of sequences, strings, and maps, the bound parameter may be omitted. If it is omitted, the bound is not specified; such a collection is referred to as “unbounded.” (All arrays must be bounded.) In that case, the type may have *no* upper bound—meaning that the collection may contain any number of elements—*or* it may have an *implicit* upper bound imposed by a given type representation (which might, for example, provide only a certain number of bits in which to store the bound) or implementation (which might, for example, impose a smaller default bound than the maximum allowed by the type representation for resource management purposes). Because of this ambiguity, type designers are encouraged to choose an explicit upper bound whenever possible.

In the cases of sequences, strings, and maps, the bound is a single value. Arrays have independent bounds on each of their dimensions; they can also be said to have an overall bound, which is the product of all of their dimensions’ bounds.

For example, a one-dimensional array of 10 integers, a one-dimensional array of 10 short integers, a sequence of at most 10 integers, and a sequence of an unspecified number of integers are all of different types. However, all one-dimensional arrays of 10 integers are of the same type.

7.2.2.4.4 Aggregated Types

Aggregations are containers for elements—“members”—of (potentially) heterogeneous types. Each member is identified by a string name and an integer ID. Each must be unique within a given type. Each member also has a type; this type may be the same as or different than the types of other members of the same aggregated type.

The relative order in which an aggregated type’s members are defined is significant, and may be relied upon by certain Data Representations.

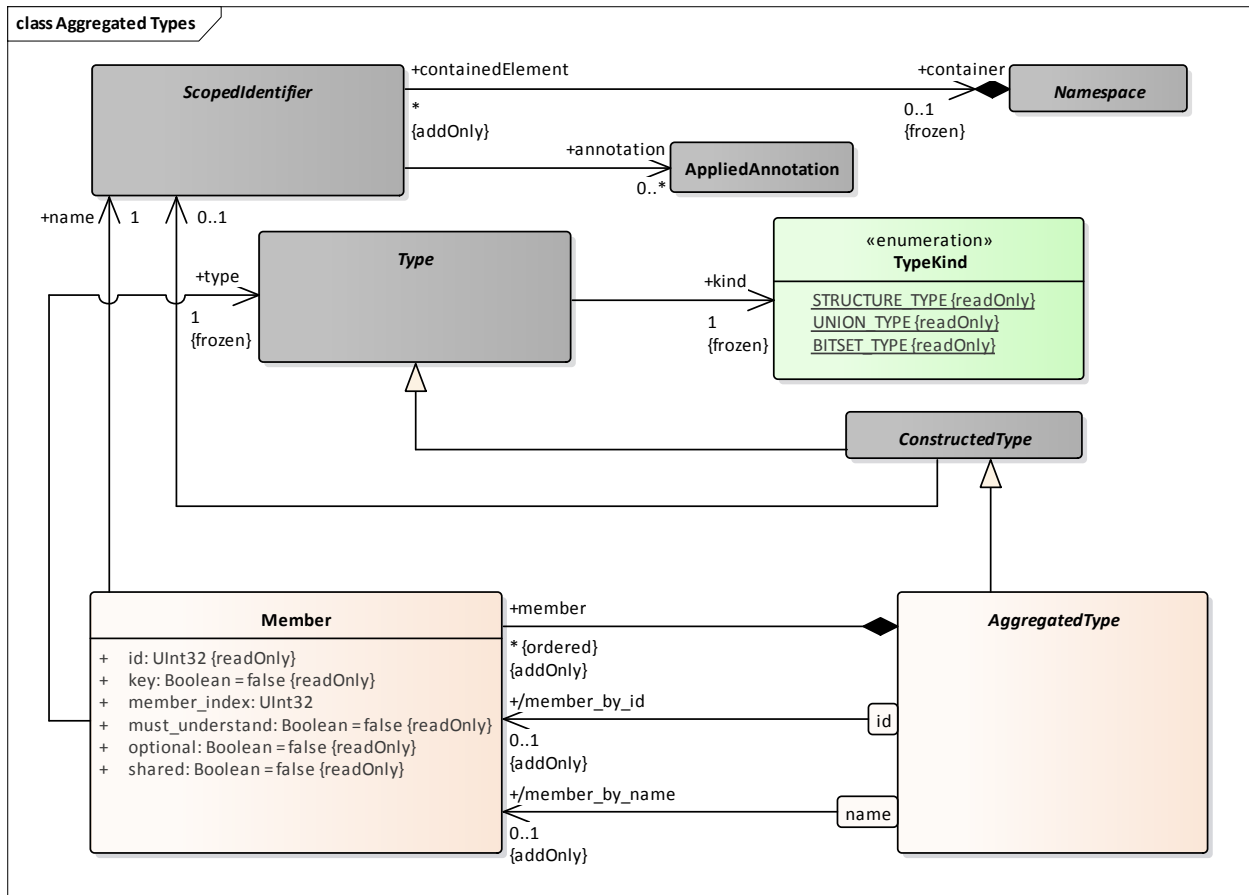


Figure 16 – Aggregated Types

There are three kinds of Aggregated Types: structures, unions, and annotations. These kinds are described in Table 8.

Table 8 – Aggregated Types

<i>Type Kind</i>	<i>Type Name</i>	<i>Description</i>
UNION_TYPE	<i>Assigned when type is defined</i>	Discriminated exclusive aggregation of members. Unions define a well-known discriminator member and a set of type-specific members.
STRUCTURE_TYPE	<i>Assigned when type is defined</i>	Non-exclusive aggregation of members. A type designer may declare any number of members within a structure. Unlike in a union, there are no implicit members in a structure, and values for multiple members may coexist.

7.2.2.4.4.1 Structure Types

A type designer may declare any number of members within a structure. Unlike in a union, there are no implicit members in a structure, and values for multiple members may coexist.

A structure can optionally extend one other structure, its “base_type.” In the event that there is a name or ID collision between a structure and its base type, the definition of the derived structure is ill-formed.

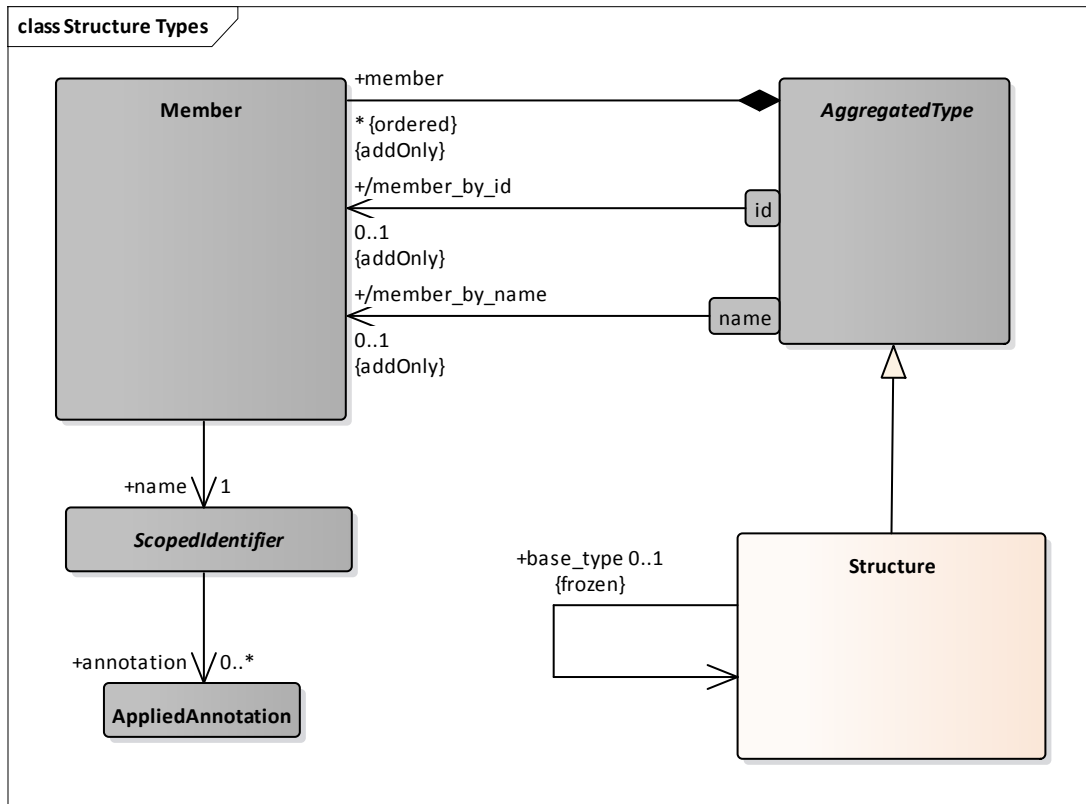


Figure 17 – Structure Types

7.2.2.4.4.2 Union Types

Unions define a well-known discriminator member and a set of type-specific members. The name of the discriminator member is always “discriminator”; that name is reserved for union types and is not permitted for type-specific union members. The discriminator member is always considered to be the first member of a union.

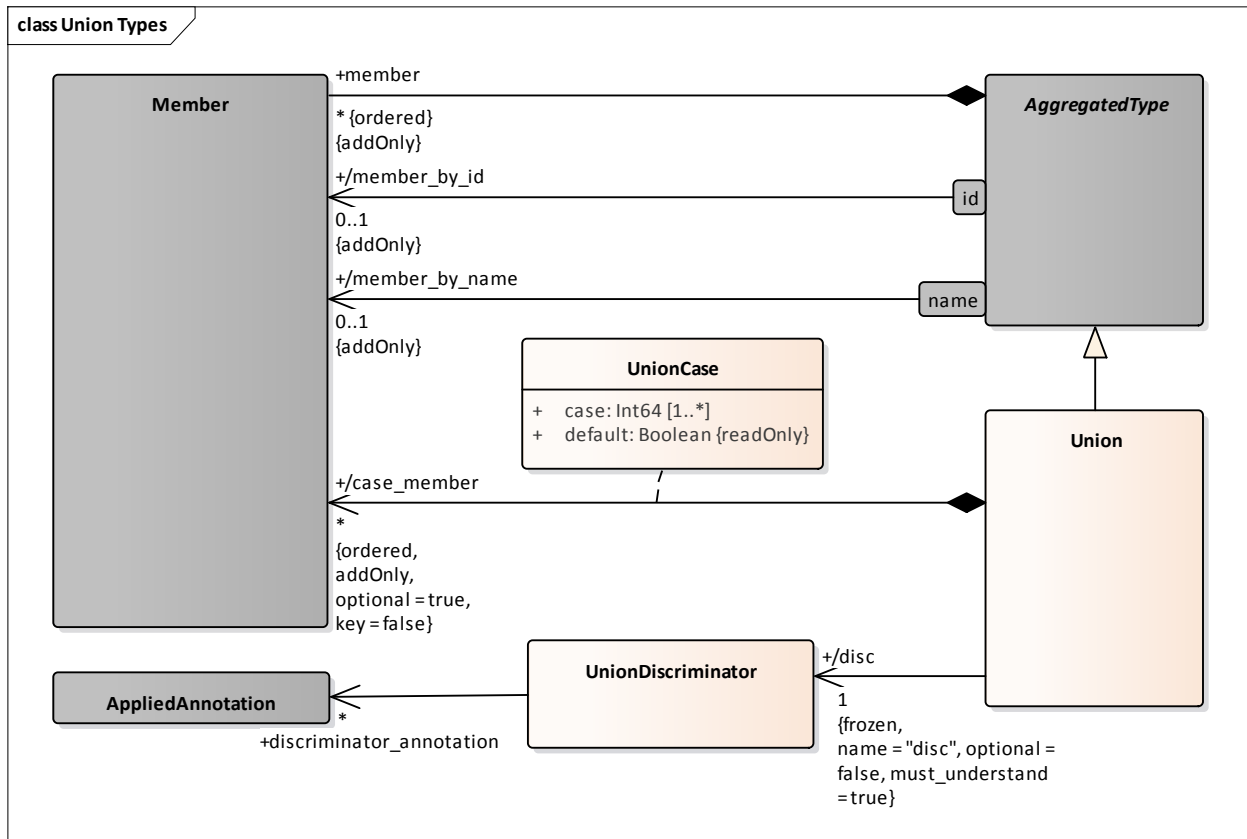


Figure 18 – Union Types

Each type-specific member is associated with one or more values of the discriminator. These values are identified in one of two ways: (1) They may be identified explicitly; it is not allowed for multiple members to explicitly identify the same discriminator value. (2) At most one member of the union may be identified as the “default” member; any discriminator value that does not explicitly identify another member is considered to identify the default member. These two mechanisms together guarantee that any given discriminator value identifies at most one member of the union. (*Note* that it is not required for every potential discriminator value to be associated with a member.) These mappings from discriminator values to members are defined by a union *type* and do not differ from object to object.

The value of the member associated with the current value of the discriminator is the only member value considered to exist in a given object of a union type at a given moment in time. However, the value of the discriminator field may change over the lifetime of a given object, thereby changing which union member’s value is observed. When such a change occurs, the initial value of the newly observed member is undefined by the type system (though it may be defined by a particular language binding). In particular, it is not defined whether, upon switching from a discriminator value *x* to a different value *y* and then immediately back to *x*, the previous value of the *x* member will be preserved.

The discriminator of a union must be of one of the following types:

- Boolean
- Byte

- Char8, Char16
- Int16, UInt16, Int32, UInt32, Int64, UInt64
- Any enumerated type
- Any alias type that resolves, directly or indirectly, to one of the aforementioned types.

7.2.2.4.4.3 Member IDs

As noted above, each member of an aggregated type is uniquely identified within that type by an integer “member ID.” Member IDs are unsigned and have a range that can be represented in 28 bits: from zero to 268,435,455 (0x0FFFFFFF). (The full range of a 32-bit unsigned integer is *not* used in order to allow binary Data Representations the freedom to embed a small amount of meta-data into a single 32-bit field if they so desire.)

The upper end of the range, from 268,419,072 (0x0FFFC000) to 268,435,455 (0x0FFFFFFF) inclusive, is reserved for use by the OMG, either by this specification—including future versions of it—or by future related specifications (16,384 values). The largest value in this range—0x0FFFFFFF—shall be used as a sentinel to indicate an invalid member ID. This sentinel is referred to by the name `MEMBER_ID_INVALID`.

The remaining part of the member ID range—from 0 to 268,402,687 (0x0FFFBFFF)—is available for use by application-defined types compliant with this specification.

7.2.2.4.4.4 Members That Must Be Understood by Consumers

A consumer of data may not have the same definition for a type as did the producer of that data. Such a situation may come about as a result of the independent, decoupled definition of the respective types or as a result of a single type’s evolution over time. A consumer, upon observing a member value it does not understand, must be able to determine whether it is acceptable to ignore the member and continue processing other members, or whether the entire data sample must be discarded.

Each member of an aggregated type has a Boolean attribute “must understand” that satisfies this requirement. If the attribute is true, a data consumer, upon identifying a member it does not recognize, must discard the entire data sample to which the member belongs. If the attribute is false, the consumer is permitted to process the sample, omitting the value of the unrecognized member.

In a structure type, each member may have the “must understand” attribute set to true or false independently.

In a union type, the discriminator member shall always have the “must understand” attribute set to true.

The ability of a consumer to detect the presence of an unrecognized member depends on the Data Representation. Each representation shall therefore define the means by which such detection occurs.

7.2.2.4.4.5 Optional Members

Each member of an aggregated type has a Boolean attribute that indicates whether it is *optional*. Every object of a given type shall be considered to contain a value for every non-optional member defined by that type. In the event that no explicit value for such a member is ever provided in a Data Representation of that object, that member is considered to nevertheless have the default “zero” value defined in **Table 9** below.

Table 9 – Default values for non-optional members

<i>Type Kind</i>	<i>Default Value</i>
BYTE	0x00
BOOLEAN	False
INT_16_TYPE, UINT_16_TYPE, INT_32_TYPE, UINT_32_TYPE, INT_64_TYPE, UINT_64_TYPE, FLOAT_32_TYPE, FLOAT_64_TYPE, FLOAT_128_TYPE	0
CHAR_8_TYPE, CHAR_16_TYPE	‘\0’
STRING_TYPE	“”
ARRAY_TYPE	<i>An array of the same dimensions and same element type whose elements take the default value for their corresponding type.</i>
ALIAS_TYPE	<i>The default type of the alias’s base type.</i>
SEQUENCE_TYPE	<i>A zero-length sequence of the same element type.</i>
MAP_TYPE	<i>An empty map of the same element type.</i>
ENUM_TYPE	<i>The first value in the enumeration.</i>
UNION_TYPE	<i>A union with the discriminator set to select the default element, if one is defined, or otherwise to the lowest value associated with any member. The value of that member set to the default value for its corresponding type.</i>
STRUCTURE_TYPE	<i>A structure without any of the optional members and with other members set to their default values based on their corresponding types.</i>

An object may omit a value for any optional member(s) defined by its type. Omitting a value is semantically similar to assigning a null value to a pointer in a programming language: it

indicates that no value exists or is relevant. Implementations shall *not* provide a default value in such a case.

Union members, including the discriminator, shall never be optional.

Structure members may be optional. The designer of a structure can choose which members are optional on a member-by-member basis.

The value of a member’s “optional” attribute is unrelated to the value of its “must understand” attribute. For example, it is legal to define a type in which a non-optional member can be safely skipped or one in which an optional member, if present and not understood, must lead to the entire sample being discarded.

7.2.2.4.4.6 Key Members

A given member of an aggregated type may be designated as part of that type’s *key*. The type’s key will become the key of any DDS Topic that is constructed using the aforementioned aggregated type as the Topic’s type. If a given type has no members designated as key members, then the type—and any DDS Topic that is constructed using it as its type it—has no key.

Key members shall never be optional, and they shall always have their “must understand” attribute set to true.

A type’s key can only include members of the following types: primitive, aggregation, enumeration, bitmask, array, and sequence. Aliases to one of the previous types can also be used as key members. Members of type map cannot be included as part of the key.

Which members may together constitute a type’s key depends on that type’s kind.

In a structure type, the key designation can be applied to any member and to any number of members.

In a union type, only the discriminator is permitted to be a key member. The union discriminator is marked as a key by annotating the discriminator itself with the `@key` annotation as shown in the example below:

```
enum CommandKind {
    START,
    STOP,
    GO_LEFT,
    GO_RIGHT
};

union MyCommand switch (@key CommandKind) {
case START:
    float delay; /* delay until start in seconds */
case STOP:
    float distance; /* distance to stop in meters */
```

```

case GO_LEFT:
case GO_RIGHT:
    float angle; /* Angle to change direction in radians */
};

```

If a member of type array or sequence is marked as a key member of an aggregated type *T*, all the elements in the array or sequence shall be considered part of the key of *T*. In the case of a sequence, the length of the sequence is also considered as part of the key ahead of the sequence elements.

In the event that the type *K* of a key member of a given type *T* itself defines key members, only the key of *K*, and not any other of its members, shall be considered part of the key of *T*. This relationship is recursive: the key members of *K* may themselves have nested key members.

For example, suppose the key of a medical record is a structure describing the individual whose record it is. Suppose also that the nested structure (the one describing the individual) has a key member that is the social security number of that individual. The key of the medical record is therefore the social security number of the person whose medical record it is.

7.2.2.4.5 Verbatim Text

System developers frequently require the ability to inject their own text into the code produced by a Type Representation compiler. Such output typically depends on the target programming language, not on the Type Representation. Furthermore, it is desirable to be able to preserve information about such output across translations of the Type Representation. Therefore, it is appropriate to manage user-specified content within the Type System for use by all Type Representations and therefore by Type Representation compilers. The `VerbatimText` class serves this purpose; each constructed type may refer to one or more instances of this class.

A `VerbatimText` object defines three properties; each is a string:

- `language`: The target programming language for which the output text applies.
- `placement`: The location within the generated output at which the output text should be inserted.
- `text`: The literal output text to be copied into the output by the Type Representation compiler.

7.2.2.4.5.1 Property: Language

When a Type Representation compiler generates code for the programming language named (case-insensitively) by this property, it shall copy the string contained in the `text` property into its output.

- The string “c” shall indicate the C programming language [C-LANG].
- The string “c++” shall indicate the C++ programming language [C++-LANG].
- The string “java” shall indicate the Java programming language [JAVA-LANG].

- The string “*” (an asterisk) shall indicate that `text` applies to all programming languages.

7.2.2.4.5.2 Property: Placement

This string identifies where, relative to its other output, the Type Representation compiler shall copy the `text` string. It shall be interpreted in a case-insensitive manner. All Type Representation compilers shall recognize the following placement strings; individual compiler implementations may recognize others in addition.

- `begin-declaration-file`: The `text` string shall be copied at the beginning of the file containing the declaration of the associated type before any type declarations.

For example, a system implementer may use such a `VerbatimText` instance to inject import statements into Java output that are required by literal code inserted by other `VerbatimText` instances.

- `before-declaration`: The `text` string shall be copied immediately before the declaration of the associated type.

For example, a system implementer may use such a `VerbatimText` instance to inject documentation comments into the output.

- `begin-declaration`: The `text` string shall be copied into the body of the declaration of the associated type before any members or constants.

For example, a system implementer may use such a `VerbatimText` instance to inject additional declarations or implementation into the output.

- `end-declaration`: The `text` string shall be copied into the body of the declaration of the associated type after all members or constants.
- `after-declaration`: The `text` string shall be copied immediately after the declaration of the associated type.
- `end-declaration-file`: The `text` string shall be copied at the end of the file containing the declaration of the associated type after all type declarations.

7.2.2.4.5.3 Property: Text

The Type Representation compiler shall copy the string contained in this property into its output as described above.

7.2.2.4.6 External Data

In some cases, it is necessary and/or desirable to provide information to a language binding that a certain member’s data should be stored, not inline within its containing type, but external to it (e.g., using a pointer).

- For example, the data may be very large, such that it is impractical to copy it into a sample object before sending it on the network. Instead, it is desirable to manage the storage outside of the middleware and assign a reference in the sample object to this external storage.

- For example, the type of the member may be the type of a containing type (directly or indirectly). This will be the case when defining linked lists or any of a number of more complex data structures.

Type Representations shall therefore allow the following type relationships in the case of external members, which would typically cause errors in the case of non-external members:

- An external member of an aggregated type shall be permitted to refer to a type whose definition is incomplete (*i.e.* is identified only by a forward declaration) at the time of the member's declaration.
- An external member of an aggregated type shall be permitted to refer to the member's containing type.

Each member of an aggregated type—with the exception of the discriminator of a union type—may be optionally marked as *external*. Likewise, the elements of a collection type may be optionally marked as external.

Note that this attribute does *not* provide a means for modeling object graphs.

7.2.2.5 Nested Types

Not every type in a user's application will be used to type DDS Topics; some types appear only as the types of members within other types. It is desirable to distinguish these two cases for the same of efficiency; for example, an IDL compiler need not generate typed `DataWriter`, `DataReader`, and `TypeSupport` classes for types that are not intended to type topics. Types that are not intended to describe topic data are referred to as *nested* types.

7.2.2.6 Annotations

An annotation describes a piece of metadata attached to a type or an element/member/literal of an aggregated/collection/enumerated type. Annotations can also be attached to the `related_type` of an alias type. An `AnnotationType` defines the structure of the metadata as a set of `AnnotationParameters` that can be assigned values when the annotation is applied. The `AnnotationParameters` are given values when the annotation is applied to an element of that other type.

The definition of an `AnnotationType` can specify the default value of each `AnnotationParameter`. `AnnotationParameters` are restricted to certain types. This allows the compiler of a `TypeRepresentation` to be able to efficiently interpret an annotation instantiation; it also simplifies expressing the parameter values as object literals in a variety of `TypeRepresentations`.

The types permitted for an `AnnotationParameter` are:

- Primitive types
- String types of `Char8` or `Char16` elements
- Enumerated types

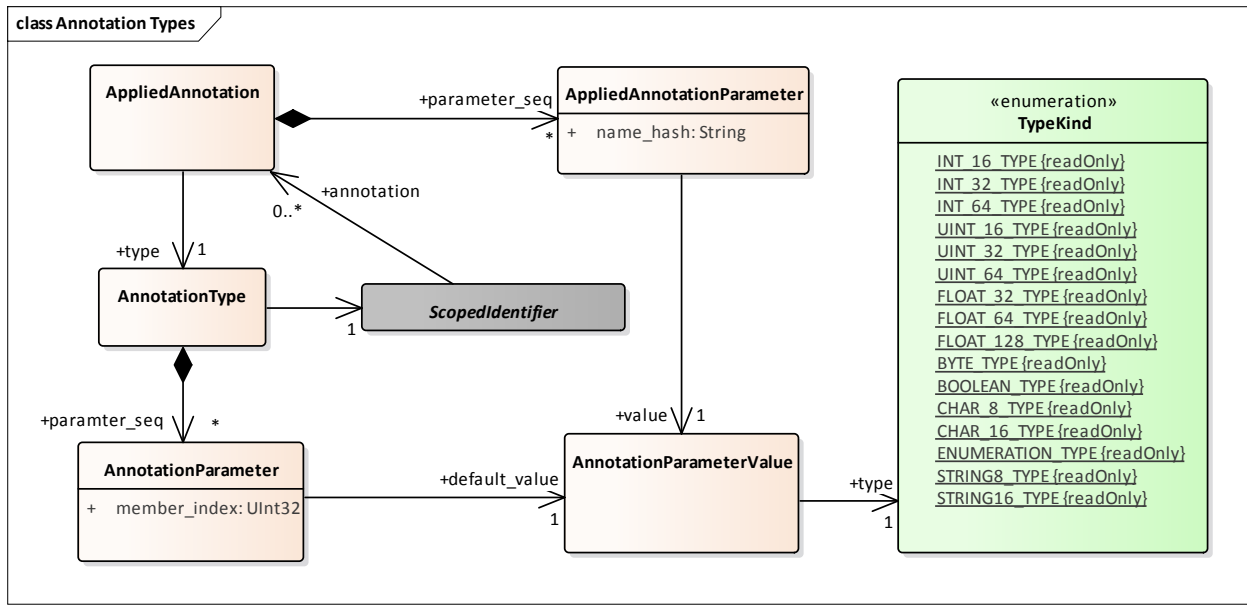


Figure 19 – Annotation Types

7.2.2.7 Try Construct behavior

Type evolution can result in a DDS `DataReader` built using type “T1” to be matched with a `DataWriter` built using a different but compatible version of the type “T2”. When the `DataReader` receives an object O2 sent by the `DataWriter` it needs to construct some object of type T1 to hold the data in O2. The expectation is that the constructed object “O1” of type T1 will faithfully capture all the information from O2 that is relevant to the application that was expecting to read objects of type T1.

There are situations where no “obviously reasonable” object of type T1 can be constructed to hold the value of a specific object “O2” of type “T2”. A type system could declare types T1 and T2 where this situation may occur to be “incompatible” thus ensuring the situation is never encountered when a `DataWriter` sends data to a matching `DataReader`. However doing so would be too restrictive for the kinds of distributed systems where DDS is deployed.

For example, a system may be deployed with `DataReader` entities reading an Aggregated type (e.g. a structure) called “STRUCT1024” with a member of type string with a maximum length of 1024 characters, see Table 10. Once the system is deployed new applications are added and the deployment extends to resource-constrained environments where the 1024 character strings can be problematic. Moreover as it turned out the value of 1024 was overly generous and in the deployed system the strings never exceed 80 characters. In this situation it becomes desirable to re-define the type as “STRUCT128”. STRUCT128 differs from STRUCT1024 in that the string member has maximum length 128, see Table 10. With these definitions there exist objects of type STRUCT1024 that cannot construct any object of type STRUCT128, namely those objects a string member of length greater than 128 characters. This is true even if the application never uses these objects. If the existence of such objects would prevent STRUCT128 from being compatible with STRUCT1024 we would not be able to adjust the type without modifying the already deployed systems, which may not be feasible.

Similar situations can occur for Collection types. For example a type “SEQ1024” that is defined as a sequence whose elements have type string with maximum length 1024 and an evolution of that type “SEQ128” that differs from SEQ1024 in that the element type is string with maximum length 128, see Table 10. Similar to the structure examples there exist be objects of type SEQ1024 that cannot construct any object of type SEQ128 and yet in many cases we do not want to consider these types as incompatible.

Table 10 – TryConstruct examples

<i>Example Type IDL definition</i>	<i>Explanation of the Type</i>
<pre>struct STRUCT1024 { string<1024> member; };</pre>	Structure Aggregated type with a member of type string with maximum length 1024 characters.
<pre>struct STRUCT128 { string<128> member; };</pre>	Structure Aggregated type with a member of type string with maximum length 128 characters.
<pre>typedef sequence< string<1024> > SEQ1024;</pre>	Sequence Collection type with element of type string with maximum length 1024 characters.
<pre>typedef sequence< string<128> > SEQ128;</pre>	Sequence Collection type with element of type string with maximum length 128 characters.

To avoid the situation described above the type compatibility relationship defined by this type system (see Clause 7.2.4) does not require that all objects of a type “T2” can faithfully construct some object of some other type “T1”, as a pre-requisite for compatibility. The type system only requires that a reasonable subset of T2 object can construct some object of type T1 and that the situations where this is not possible are detected and gracefully handled. The rules for this are formally defined in Clause 7.2.4.

Therefore even when two types T1 and T2 are compatible it may be possible to encounter an object sent by a `DataWriter` of type T2 that cannot be used to construct any object of the T1 type expected by the `DataReader` without losing some potentially critical information. For example, depending on the application truncating a 20-character string sent by the `DataWriter` into a 10-character string that may be the maximum allowed by the `DataReader` could result in misinterpretation and application malfunction. The same could be said for trimming a received sequence to a shorter length.

If no “reasonable” T1 object can be constructed from a given object O2 of type T2, we say that **“O2 cannot construct any object of type T1”**.

Object construction for collection and aggregated types is done recursively. To construct the collection/aggregated object it is necessary to construct all nested elements/members. For this reason failure to construct a nested element/member can prevent the construction of the collection/aggregated type.

There are situations when it is not desirable to fail the construction of a collection or aggregated object of type T1 just because some nested element/member cannot be constructed. The failure to construct the element/member would cause all other nested elements/members to be “lost” and not just the problematic one. In some situations it may be more desirable to trim the problematic member or set it to some well-known default value. To support these scenarios Collection and Aggregated types may explicitly declare the `TryConstruct` behavior of each of their elements or members.

- Array and Sequence collection types may explicitly declare that their element has one of three kinds of `TryConstruct` behavior, see Table 11.
- Map collection types may explicitly declare that their “key” and or “value” element has one of three kinds of `TryConstruct` behavior, see Table 11.
- Structure and Union types may explicitly declare member has one of three kinds of `TryConstruct` behavior, see Table 11. In the case of Unions this extends to the discriminator member.

The `TryConstruct` behavior kinds are described in Table 11 below. The default behavior unless otherwise specified using the `TryConstruct` annotation is `DISCARD`.

Table 11 – TryConstruct behavior kinds

<i>TryConstruct kind</i>	<i>Description</i>
DISCARD	<p>Failure to construct an element or member propagates to the collection or aggregated type that contains it.</p> <p>If an element or member cannot be constructed, then the collection or aggregated object that contains the element or member cannot be constructed either.</p>
USE_DEFAULT	<p>Failure to construct an element or member is contained—element or member is set to its default value.</p> <p>If an element or member cannot be constructed, the element/member shall be set to its default value (according to its type as described in Table 9) and does not cause the collection/aggregated object to fail its construction.</p>

TRIM	<p>Failure to construct an element or member is contained—element or member is trimmed.</p> <p>This option only applies to elements or members of type string, wide string, sequence, or map. The behavior when applied to other element/member types is unspecified and may be treated as an error.</p> <p>The option affects the situation where failure to construct is due to the length of the collection sent exceeding what can be accommodated on the receiving member collection type.</p> <p>In this situation the element or member is constructed trimming the received object to the length that can be accommodated by the receiving member type. The order of the characters in the string or elements in the sequence or map is preserved.</p>
------	---

7.2.3 Type Extensibility and Mutability

In some cases, it is desirable for types to evolve without breaking interoperability with deployed components already using those types. For example:

- A new set of applications to be integrated into an existing system may want to introduce additional fields into a structure. These new fields can be safely ignored by already deployed applications, but applications that do understand the new fields can benefit from their presence.
- A new set of applications to be integrated into an existing system may want to increase the maximum size of some sequence or string in a Type. Existing applications can receive data samples from these new applications as long as the actual number of elements (or length of the strings) in the received data sample does not exceed what the receiving applications expects. If a received data sample exceeds the limits expected by the receiving application, then the sample can be safely ignored (filtered out) by the receiver.

In order to support use cases such as these, the type system introduces the concept of *appendable* and *mutable* types.

- A type may be *FINAL*, indicating that the range of its possible data values is strictly defined. In particular, it is not possible to add elements to members of collection or aggregated types while maintaining type assignability.
- A type may be *APPENDABLE*, indicating that two types, where one contains all of the elements/members of the other plus additional elements/members appended to the end, may remain assignable. Note that this was called *EXTENSIBLE* in xtypes version 1.1 and prior.
- A type may be *MUTABLE*, indicating that two types may differ from one another in the additional, removal, and/or transposition of elements/members while remaining assignable.

This attribute may be used by the Data Representations to modify the encoding of the type in order to support its extensibility.

The meaning of these extensibility kinds is formally defined with respect to type compatibility in Clause 7.2.4, “Type Compatibility.” It is summarized more generally in Table 12.

Table 12 – Meaning of marking types as appendable

<i>Type Kind</i>	<i>Meaning of marking type as appendable</i>
Aggregated Types: STRUCTURE_TYPE, UNION_TYPE	Aggregated types may be final, appendable, or mutable on a type-by-type basis. However, the extensibility kind of a structure type with a base type must match that of the base type. It shall not be permitted for a subtype to change the extensibility kind of its base type. Any members marked as keys must be present in all variants of the type.
Collection Types: ARRAY_TYPE, SEQUENCE_TYPE, MAP_TYPE	Sequence and map types are always mutable. Array types are always final. Variations of a mutable collection type may change the maximum number of elements in the collection.
Enumerated Types: ENUMERATION_TYPE, BITMASK_TYPE	Enumerated types may be final, appendable, or mutable on a type-by-type basis. Bitmask types are always final.
String Types: STRING8_TYPE, STRING16_TYPE	String types are always mutable.
ALIAS_TYPE	Since aliases are semantically equivalent to their base types, the extensibility kind of an alias is always equal to that of its base type.
Primitive types	Primitive types are always final.

7.2.4 Type Compatibility

In order to maintain the loose coupling between data producers and consumers, especially as systems change over time, it is desirable that the two be permitted to use slightly different versions of a type, and that the infrastructure perform any necessary translation. To support type evolution and inheritance the type system defines the “is-assignable-from” directed binary relationship between every pair of types in the Type System.

Given two types T1 and T2, we will write:

T1 is-assignable-from T2

...if and only T1 is related to T2 by this relationship. The rules to determine whether two types have this relationship are given in the following subclauses.

Intuitively, if T1 *is-assignable-from* T2, it means that in general it is possible, in a structured way, to set the contents of an object of type T1 to the contents of an object of T2 (or perhaps a subset of those contents, as defined below) without leading to incorrect interpretations of that information.

7.2.4.1 Constructing objects of one type from objects of another type

The fact that T1 *is-assignable-from* T2, does not mean that *all* objects of T2 can be used to construct an object of type T1 (for example, a collection may have too many elements).

What the *is-assignable-from* indicates is that the difference between T2 and T1 is such that (a) a meaningful subset of T2 objects can construct T1 objects without misinterpretation and that (b) the remaining objects of T2—which cannot construct T1 objects—can be detected as such so that misinterpretations can be prevented. For the sake of run-time efficiency, these per-object “can-construct” rules are designed such that their enforcement does not require any inspection of a data producer’s type definition. Per-object enforcement can potentially be avoided altogether—depending on the implementation—by declaring a type to be *final*², forcing producer and consumer types to match exactly; see Clause 7.2.3.

In the case T1 *is-assignable-from* T2 but an object O2 of type T2 is encountered that cannot construct any object of type T1, the default behavior is to discard the O2 object to avoid misinterpretation. This behavior can be altered when the object O2 is a member of an Aggregated type (e.g. a structure). In this case the behavior is determined by the TryConstruct behavior specified for the member. See Clause 7.2.2.7.

Therefore, for each pair of types T1 and T2 this specification defines the rules for T1 to be assignable-from T2. Assuming T1 is-assignable-from T2 the specification also defines which objects of type T2 can be used to construct an object of type T1.

For example:

Table 13 – Type assignability example

<i>T1</i>	<i>T2</i>	<i>Type compatibility</i>	<i>Object construction</i>
Sequence of 10 integers	Sequence of 5 integers	<p>T1 is assignable from T2: All objects of type T2 can be used to initialize T1 objects.</p> <p>T2 is assignable from T1: All objects of type T1 can either be used to construct an object of type T1 or reliably detected that that cannot initialize T1.</p>	<p>Any object O2 of type T2 can construct an object of type T1.</p> <p>Only T1 objects containing at most 5 elements can construct T2 objects.</p>

² DDS-based systems have an additional tool to enforce stricter static type consistency enforcement: the TypeConsistencyEnforcementQosPolicy. See Clause 7.6.2.3.

7.2.4.2 Concept of Delimited Types

Delimited types are those types “T” whose serialized object representation is such that the receivers of an object of that type “T” who only know a type T1 assignable-from type “T” are able to reliably delimit the object within the serialized representation. This means that where appropriate the receiver may “skip” that object and proceed to process other objects that are serialized after.

Primitive and Enumerated types (Enumeration and Bitmask) are delimited types as their serialized size is fixed.

Strings and wide strings are delimited types because the serialization starts with a size from which it is possible to derive the overall serialized length of the string.

Collection types (arrays, sequences, maps) are delimited if the collection element type is delimited. In the case of a map collection the key type must also be delimited. Otherwise the collection is not delimited. The reason is that the receiver of a compatible collection type always knows the length of the collection: Either it is encoded in the serialized representation (sequences and maps) or it is the same as the receiver type in the case of arrays.

Other than the types mentioned above all other types with extensibility kind FINAL **are not** delimited.

Types with extensibility kind APPENDABLE **are** delimited if serialized with encoding version 2 (DELIMITED_CDR). See Clause 7.4.2. They **are not** delimited if serialized with encoding version 1.

Mutable types are also delimited with both encoding version 1 and encoding version 2.

- The serialized representation used for version 1 encoding (PL_CDR) is a list of length-encoded elements ended by a sentinel, which delimits the serialized object. See Clause 7.4.1.2.
- The serialized representation used for version 2 encoding (PL_CDR2) starts with a delimiter header similar to the one used for DELIMITED_CDR, which delimits the serialized object.

7.2.4.3 Strong Assignability

If types T1 and T2 are equivalent using the MINIMAL relation (see Clause 7.3.4.7), or alternatively if T1 is assignable-from T2 and T2 is a delimited type, then T1 is said to be “strongly” assignable from T2.

7.2.4.4 Assignability Rules

7.2.4.4.1 Assignability of Equivalent Types

If two types T1 and T2 are equivalent according to the MINIMAL relation (see Clause 7.3.4.7), then they are mutually assignable, that is, T1 is-assignable-from T2 and T2 is-assignable-from T1.

The reverse is not always true. The type system contains mutually assignable types that are not equivalent according to the MINIMAL relation.

7.2.4.4.2 Non-serialized Members

Members that are marked as non-serialized, see Sub Clause 7.3.1.2.1.13, shall be ignored during type compatibility checking.

7.2.4.4.3 Alias Types

Table 14 – Definition of the *is-assignable-from* relationship for alias types

<i>T1 Type Kind</i>	<i>Type assignability</i>	<i>Object construction</i>
ALIAS_TYPE	Any non ALIAS_TYPE type kind T2 if and only if T1.base_type is-assignable-from T2	Construct according to the rules for constructing T1.base_type objects from T2 objects
Any non ALIAS_TYPE type kind	ALIAS_TYPE T2 if and only if T1 is-assignable-from T2.base_type	Construct T1 objects according to the rules for constructing T1 from objects of type T2.base_type
ALIAS_TYPE	ALIAS_TYPE if and only if T1.base_type is-assignable-from T2.base_type	Construct according to the rules for constructing T1.base_type objects from T2.base_type objects

For the purpose of evaluating the *is-assignable-from* relationship, aliases are considered to be fully resolved to their ultimate base types. For this reason, alias types are not discussed explicitly in the subsequent clauses. Instead, if T is an alias type, then it shall be treated as if T == T.base_type.

7.2.4.4.4 Primitive Types

Table 15 below defines the *is-assignable-from* relationship for Primitive Types. These conversions are designed to preserve the data during translation. Furthermore, in order to preserve high performance, they are designed to enable the preservation of *data representation*, such that a `DataReader` is not required to parse incoming samples differently based on the `DataWriter` from which they originate. (For example, although a short integer could be promoted to a long integer without destroying information, a binary Data Representation is likely to use different amounts of space to represent these two data types. If, upon receiving each sample from the network, a `DataReader` does not consult the type definition of the `DataWriter` that sent that sample, it would not know how many bytes to read. The runtime expense of this kind of type introspection on the critical path is undesirable.)

Table 15 – Definition of the *is-assignable-from* relationship for primitive types

<i>T1 Type Kind</i>	<i>T2 Type Kinds for which T1 is-assignable-from T2 Is True</i>	<i>Object construction</i>
Any Primitive Type	The same Primitive Type	Copy the primitive object.

BYTE_TYPE	BITMASK_TYPE if and only if T2.bound is between 1 and 8, inclusive.	For each bitflag that is set in the bitmask construct the integer value (1 << position) using the position of that bitflag. Add all those integer values to obtain the resulting object O1 of type T1
UINT16_TYPE	BITMASK_TYPE if and only if T2.bound is between 9 and 16, inclusive.	
UINT32_TYPE	BITMASK_TYPE if and only if T2.bound is between 17 and 32, inclusive.	
UINT64_TYPE	BITMASK_TYPE if and only if T2.bound is between 33 and 64, inclusive.	

7.2.4.4.5 String Types

The *is-assignable-from* relationship for string types is described in Table 16.

Table 16 – Definition of the is-assignable-from relationship for string types

<i>T1 Type Kind</i>	<i>T2 Type Kinds for which T1 is-assignable-from T2 Is True</i>	<i>Object construction (assuming type assignability)</i>
STRING_TYPE	STRING_TYPE if and only if T1.element_type <i>is-assignable-from</i> T2.element_type	An object O2 of type T2 can-construct an object of type T1 if and only if O2.length <= T1.length Copy each character. O1.length is set to O2.length.

7.2.4.4.5.1 Example: Strings

According to the above rules, any string type of narrow characters is assignable from any other string type of narrow characters. Any string type of wide characters is assignable from any other string type of wide characters. However, string types of narrow characters are not assignable from string types of wide characters, because of the possibility of data misinterpretation. For example, suppose a string of wide characters is encoded using the CDR Representation. If a consumer of strings of narrow characters were to attempt to consume that string, it might read consider the first byte of the first character to be a character onto itself, the second byte of the first character to be a second character, and so on. The result would be a string of narrow characters having “junk” contents.

Furthermore, any T2 string *object* containing more characters than the bound of the T1 string type cannot construct any object of type T1 in order to prevent data misinterpretations resulting from truncations. For example, consider two versions of a shopping list application. The list of purchases is represented by a sequence of strings. Version 2.0 of the application increased the bounds of these strings. Supposing that the list items “cat food” and “catsup” were too long to be

understood by a version 1.0 consumer, it would be better to come home from the store without either item than to come home with two cats instead.

7.2.4.4.6 Collection Types

The *is-assignable-from* relationship for collection types is based in part on the same relationship as applied to their element types.

Table 17 – Definition of the *is-assignable-from* relationship for collection types

<i>T1 Type Kind</i>	<i>T2 Type Kinds for which T1 is-assignable-from T2 Is True</i>	<i>Object construction (assuming type assignability)</i>
ARRAY_TYPE	<p>ARRAY_TYPE if and only if³:</p> <ul style="list-style-type: none"> • T1.bounds[] == T2.bounds[] • T1.element_type is strongly assignable from T2.element_type 	<p>To construct an object of type T1 from an object O2 of type T2:</p> <p>Each element of the T1 array shall be constructed from the corresponding element of the O2 array.</p> <p>If an element of T1 cannot be constructed from the O2 element, the result depends on the TryConstruct behavior associated with T1 element type.</p> <ul style="list-style-type: none"> • If DISCARD, O2 cannot construct any object of type T1. • If USE_DEFAULT or TRIM, the element is constructed accordingly and the array of type T1 is successfully constructed.

³ Design rationale: This specification allows sequence, map, and string bounds to change but not array bounds. This is because of the desire to avoid requiring the consultation of per-DataWriter type definitions during sample deserialization. Without such consultation, a reader of a compact data representation (such as CDR) will have no way of knowing what the intended bound is. Such is not the case for other collection types, which in CDR are prefixed with their length.

SEQUENCE_TYPE	<p>SEQUENCE_TYPE if and only if T1.element_type is strongly assignable from T2.element_type</p>	<p>An object O2 of type T2 can construct T1 if and only if O2.length <= T1.length</p> <p>O1.length is set to O2.length.</p> <p>Construct each in O1 from the corresponding O2 element.</p> <p>If an element of O2 cannot construct T1.element_type, the result depends on the TryConstruct behavior associated with T1 element type.</p> <ul style="list-style-type: none"> • If DISCARD, O2 cannot construct any object of type T1. • If USE_DEFAULT or TRIM, the element is constructed accordingly and the O1 sequence is successfully constructed.
MAP_TYPE	<p>MAP_TYPE if and only if:</p> <ul style="list-style-type: none"> • T1.key_element_type is strongly assignable from T2.key_element_type • T1.element_type is strongly assignable from T2.element_type. 	<p>An object O2 of type T2 can construct T1 if and only if O2.length <= T1.length</p> <p>The constructed object O1 shall be as if the O1 map were cleared of all elements and subsequently all T2 map entries were added to it. The entries are not logically ordered.</p> <p>If a key element of O2 cannot construct the corresponding key type of T1 the entire map O2 cannot construct any object of type T1.</p> <p>If a value element of O2 cannot construct T1.element_type, the result depends on the TryConstruct behavior associated with T1 element type.</p> <ul style="list-style-type: none"> • If DISCARD, O2 cannot construct any object of type T1. • If USE_DEFAULT or TRIM, the element is constructed accordingly and the O1 object is successfully constructed.

7.2.4.4.7 Enumerated Types

Conversions of bitmask, and enumerated types are designed to preserve the data during translation.

Table 18 – Definition of the *is-assignable-from* relationship for bitmask, and enumerated types

<i>T1 Type Kind</i>	<i>T2 Type Kinds for which T1 is-assignable-from T2 Is True</i>	<i>Object construction</i>
BITMASK_TYPE	BITMASK_TYPE if and only if T1.bound == T2.bound	Preserve bit values by index for all bits identified in both T1 and T2.
	UINT_32_TYPE if and only if T1.bound is between 17 and 32, inclusive.	
	UINT_16_TYPE if and only if T1.bound is between 9 and 16, inclusive.	
	UINT_64_TYPE if and only if T1.bound is between 33 and 64, inclusive.	
	BYTE if and only if T1.bound is between 1 and 8, inclusive.	
ENUMERATION_TYPE	ENUMERATION_TYPE if an only if: <ul style="list-style-type: none"> • T1.extensibility == T2.extensibility • Any literals that have the same name in T1 and T2 also have the same value, and any literals that have the same value in T1 and T2 also have the same name. • The default literal has the same value. • If extensibility is final the set of literals should be identical. Otherwise the two types should have at least one other literal (in addition to the default one) in common. 	Choose the corresponding T1 literal if it exists. If the name or value of the T2 object does not exist in T1, the object cannot construct any object of type T1.

7.2.4.4.8 Aggregated Types

For aggregated types, *is-assignable-from* is based on the same relationship between the types' members. The correspondence between members in the two types is established based on their respective member IDs and on their respective member names.

Table 19 – Definition of the *is-assignable-from* relationship for aggregated types

<i>T1 Type Kind</i>	<i>T2 Type Kinds for which T1 is-assignable-from T2 Is True</i>	<i>Object construction</i>
UNION_TYPE	<p>UNION_TYPE if and only if it is possible to unambiguously select the appropriate T1 member based on the T2 discriminator value and to transform both the discriminator and the selected member correctly. Specifically:</p> <ul style="list-style-type: none"> • T1.extensibility == T2.extensibility. • T1.discriminator.type <i>is-strongly-assignable-from</i> T2.discriminator.type. • Either the discriminators of both T1 and T2 are keys or neither are keys. • Any members in T1 and T2 that have the same name also have the same ID and any members with the same ID also have the same name. • For all non-default labels in T2 that select some member in T1 (including selecting the member in T1's default label), the type of the selected member in T1 is assignable from the type of the T2 member. • If any non-default labels in T1 that select the default member in T2, the type of the member in T1 is assignable from the type of the T2 default member. • If T1 and T2 both have default labels, the type associated with T1 default member is assignable from the type associated with T2 default member. • If T1 (and therefore T2) extensibility is final then the set of labels are identical. Otherwise, they have at least one common label other than the default label. 	<p>A union object O2 of type T2 can construct an object of type T1 if and only if:</p> <ul style="list-style-type: none"> • Either the value of O2.discriminator can construct the type of T1's discriminator. Or else the discriminator has TryConstruct behavior set to DEFAULT. <p>AND</p> <ul style="list-style-type: none"> • Either the selected member "m2" in O2, if any, can construct the selected member "m1" of T1, if any (where m1 and/or m2 may be the default member). Or else the selected member (if any) has TryConstruct behavior set to DEFAULT or TRIM. <p>Assuming O2 can construct an object of type T1, then:</p> <ul style="list-style-type: none"> • The constructed object O1 discriminator is constructed from the object O2's discriminator or if that is not possible it is set according to its TryConstruct behavior. <p>If the discriminator value selects a member m2 in O2 (which may be the default value), then:</p>

		<ul style="list-style-type: none"> • If the discriminator value also selects a member m1 in O1 (which may be the default value), then m1 is constructed from m2 or if that is not possible it is set according to its TryConstruct behavior. • If the discriminator value does not select any member in O1, then there is no value assigned from m2 (i.e. m2 is “truncated”). <p>If the discriminator value does not select any member in O2, then:</p> <ul style="list-style-type: none"> • If the discriminator value selects a member m1 in O1, then m1 is set to its default value according to its type. • If the discriminator value does not select any member in T1, then there is nothing else to assign or set on T1.
STRUCTURE_TYPE	<p>STRUCTURE_TYPE if and only if:</p> <ul style="list-style-type: none"> • T1 and T2 have the same number of members in their respective keys. • For each member “m1” that forms part of the key of T1 (directly or indirectly), there is a corresponding member “m2” that forms part of the key of T2 (directly or indirectly) with the same member id ($m1.id == m2.id$) where <i>m1.type is-assignable-from m2.type</i>. 	<p>Each member “m1” of the T1 object takes the value of the T2 member with the same ID or name, if such a member exists.</p> <p>Each non-optional member in a T1 object that is not present in the T2 object takes its default value.</p> <p>Each optional member in a T1 object that is not present in the T2 object takes no value.</p>

	<p>(The previous two rules assure that the key of T2 can be transformed faithfully into the key of T1 without aliasing or loss of information.)</p> <ul style="list-style-type: none"> • Any members in T1 and T2 that have the same name also have the same ID and any members with the same ID also have the same name. • For each member “m1” in T1, if there is a member m2 in T2 with the same member ID, then m1.type <i>is-assignable-from</i> m2.type. • Members for which both <code>optional</code> is false and <code>must_understand</code> is true in either T1 or T2 appear in both T1 and T2. • Empty type intersections prevent assignability: There is at least one member “m1” of T1 and one corresponding member “m2” of T2 such that <code>m1.id == m2.id</code>. • <code>T1.extensibility == T2.extensibility</code> <p>AND if T1 is appendable, then any members whose member ID appears both in T1 and T2 have the same setting for the ‘optional’ attribute and the T1 member type is strongly assignable from the T2 member type.</p> <p>AND if T1 is final, then they meet the same condition as for T1 being appendable and in addition T1 and T2 have the same set of member IDs.</p> <p>For the purposes of the above conditions, members belonging to base types of T1 or T2 shall be considered “expanded” inside T1 or T2 respectively, as if they had been directly defined as part of the sub-type.</p>	<p>If a “must understand” member in the T2 object is present, then T1 must have a member with the same member ID. Otherwise the object cannot construct T1. This behavior is not affected by the TryConstruct setting.</p> <p>If a member cannot construct the corresponding member in T1, then the behavior is determined by the TryConstruct setting of the member.</p>
--	--	---

7.2.4.4.8.1 Example: Type Truncation

Consider the following type for representing two-dimensional Cartesian coordinates:

```
struct Coordinate2D {
    long x;
    long y;
};
```

(This example uses the IDL Type Representation. However, the same principles apply to any other type representation.)

Now suppose that another subsystem is to be integrated. That subsystem is capable of representing three-dimensional coordinates:

```
struct Coordinate3D {
    long x;
    long y;
    long z;
};
```

(The type `Coordinate3D` may represent a new version of the `Coordinate2D` type, or the two coordinate types may have been developed concurrently and independently. In either case, the same rules apply.)

`Coordinate2D` is assignable from `Coordinate3D`, because that subset of `Coordinate3D` that is meaningful to consumers of `Coordinate2D` can be extracted unambiguously. In this case, consumers of `Coordinate2D` will observe the two-dimensional projection of a `Coordinate3D`: they will observe the `x` and `y` members and ignore the `z` member.

7.2.4.4.8.2 Example: Type Inheritance

Type inheritance is a special case of type truncation, which allows objects of subtypes to be substituted in place of objects of supertypes in the conventional object-oriented fashion.

Consider the following type hierarchy:

```
<struct name="Vehicle">
    <member name="km_per_hour" type="int32"/>
</struct>
<struct name="LandVehicle" baseType="Vehicle">
    <member name="num_wheels" type="int32"/>
</struct>
```

(This example uses the XML Type Representation. However, the same principles apply to any other type representation.)

`LandVehicle` is assignable from `Vehicle`. Any consumer of the latter that receives an instance of the former will observe the value of the member `km_per_hour` and ignore the member `num_wheels`.

7.2.4.4.8.3 Example: Type Refactoring

As systems evolve, it is sometimes desirable to refactor data from place in a type hierarchy to another place. For example, consider the following representation of a giraffe:

```
struct Animal {
    long body_length;
    long num_legs;
};

struct Giraffe : Animal {
    long neck_length;
};
```

(This example uses the IDL Type Representation. However, the same principles apply to any other type representation.)

Now suppose that a later version of the system needs to model snakes in addition to giraffes. Snakes are also animals, but they don't have legs. We could just say that they have zero legs, but then should we add `num_scales` to `Animal` and set that to zero for giraffes? It would be better to refactor the model to capture the fact that legs are irrelevant to snakes:

```
struct Animal {
    long body_length;
};

struct Mammal : Animal {
    long num_legs;
};

struct Giraffe : Mammal {
    long neck_length;
};

struct Snake : Animal {
    long num_scales;
};
```

Because the is-assignable-from relationship is evaluated as if all member definitions were flattened into the types under evaluation, both versions of the `Giraffe` type are assignable to one another. Producers of one can communicate seamlessly with consumers of the other and correctly observe values for all fields.

7.3 Type Representation

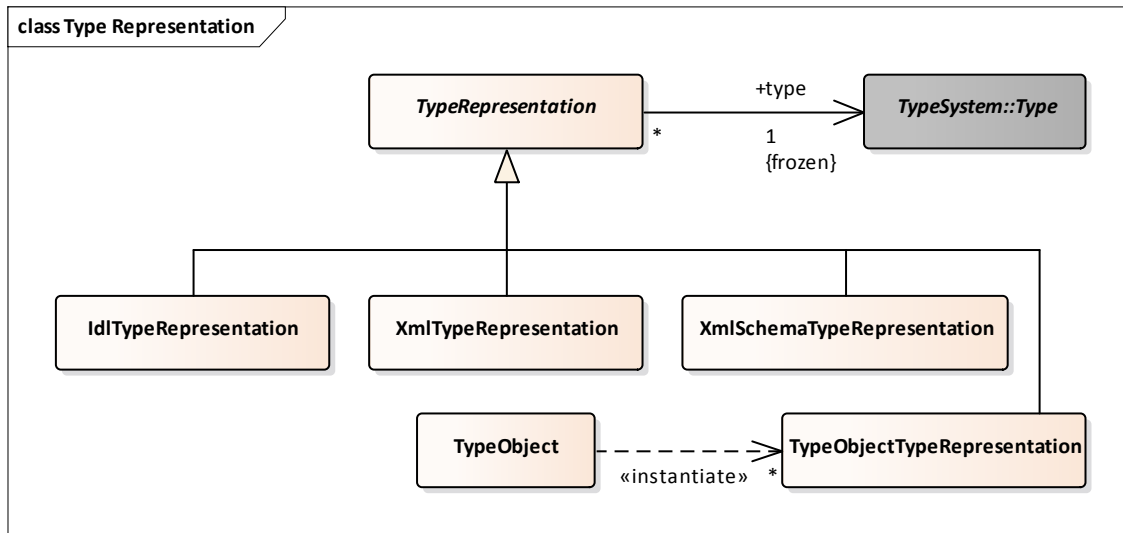


Figure 20 – Type Representation

The Type Representation module specifies the ways in which a type can be externalized so that it may be stored in a file or communicated over the network. Type Representations serve multiple purposes such as:

- Allow a user to describe and document the data type.
- Provide an input to tools that generate code and language-specific constructs to program and manipulate objects of that type.
- Provide an input to tools that want to “parse” and interpret data objects dynamically, without compile-time knowledge of the schema.
- Communicate data types via network messages so that applications can dynamically discover each other’s types or evaluate whether relationships such as *is-assignable-from* are true or false.

This specification introduces multiple equivalent Type Representations. The reason for defining multiple type representations is that each of these is better suited or optimized for a particular purpose. These representations are all equivalent because they describe the same Type System. Consequently, other than convenience or performance, there is no particular reason to use one versus the other.

The alternative representations are summarized in Table 20.

Table 20 – Alternative Type Representations

<i>Type Representation</i>	<i>Reasons for using it</i>	<i>Disadvantages</i>
IDL	<p>Compact Language. Easy to read and write by humans.</p> <p>Familiar to programmers. Uses constructs close to those in programming languages.</p> <p>Has standard language bindings to most programming languages.</p>	<p>Perceived as a legacy language by users who prefer XML-based languages.</p> <p>Not as many tools available (parsers, transformations, syntax-aware editors) as XML-languages.</p> <p>Parsing is complex.</p> <p>Requires extensions to support all concepts in the Type System, e.g. keys, optional members, map types, and member IDs.</p>
TypeObject	<p>Can provide most compact binary representation.</p> <p>Best suited for communication over a network or as an internal representation of a type.</p>	<p>Not human readable or writable.</p>
XML	<p>Compact XML language. Easy to read and write by humans.</p> <p>Defined to precisely fit the Type System so all concepts (including keys, optional member, etc.) map well.</p> <p>Syntax can be described using XSD allowing the use of editors that assist and verify the syntax of the type.</p> <p>Well-suited for run-time processing due to availability of packages that parse XML.</p>	<p>New language. Based on XML but with a schema that is previously unknown to users.</p>
XSD	<p>Popular standard. Familiar to many users. Human readable.</p> <p>Allows reusing of types defined for other purposes (e.g. web-services).</p>	<p>Cumbersome syntax. XSD was conceived as a way to define the syntax of XML documents, not as a way to define data types.</p>

	Availability of tools to do syntax checking and editors that assist with auto-completion.	No direct support for many of the constructs (e.g keys) or the types in the type model (e.g. arrays, unions, enums), resulting in having to use specific patterns that are hard to remember and error-prone. Very verbose. Hard to read by a programmer.
--	---	---

7.3.1 IDL Type Representation

The type system defined by this specification is designed to allow types to be easily represented using IDL [IDL41] with minimal extensions.

7.3.1.1 IDL Compatibility

This specification considers two aspects of IDL compatibility:

- *Backward compatibility with respect to type definitions*: Existing IDL type definitions for use with DDS remain compatible to the extent that those definitions were standards-compliant and based on implementation-independent best practices.
- *Forward compatibility with respect to IDL compilers*: With a few exceptions, IDL type definitions formulated according to this specification will be accepted by IDL compilers that do not conform to this specification.

7.3.1.1.1 Backward Compatibility with Respect to Type Definitions

This specification uses a subset of the IDL type definition syntax defined in [IDL41]. In particular, it uses the Extensible DDS Profile (Sub Clause 9.3.2 [IDL41]), which is composed of the following elements:

- Building Blocks
 - Core Data Types (Sub Clause 7.4.1 [IDL41])
 - Extended Data Types (Sub Clause 7.4.13 [IDL41])
 - Anonymous Types (Sub Clause 7.4.14 [IDL41])
 - Annotations (Sub Clause 7.4.15 [IDL41])
- Group of Annotations
 - General Purpose (Sub Clause 8.3.1 [IDL41])
 - Data Modeling (Sub Clause 8.3.2 [IDL41])
 - Data Implementation (Sub Clause 8.3.4 [IDL41])

- Code Generation (Sub Clause 8.3.5 [IDL41]).

This specification retains well-established IDL type definition syntax, such as enumeration, structure, union, and sequence definitions.

Some DDS users may be using constructs for implementation-specific purposes outside the building blocks and group of annotations listed above. These constructs remain legal for use in IDL files provided to IDL compilers compliant with this specification. However, their meanings are undefined with respect to this specification. Compilers that do not support them shall ignore them or issue a warning rather than halting with an error.

7.3.1.1.2 Forward Compatibility with Respect to Compilers

This specification retains well-established IDL type definition syntax, such as enumeration, structure, union, and sequence definitions. This degree of backward compatibility also provides forward compatibility with respect to IDL compilers.

However, this specification also defines new Type System concepts that necessarily had no defined IDL representation, such as maps and annotations. In some cases, such as with annotations, a syntax exists that does not harm compatibility; see Clause 7.3.1.2.6. In other cases, incompatibility is unavoidable.

The following pragma declarations allow IDL type designers to indicate to their tools and to human readers that their IDL file (or a portion of it) makes use of constructs defined by this specification:

```
#pragma dds_xtopics begin [<version_number>]
// IDL definitions
#pragma dds_xtopics end [<version_number>]
```

The optional version number indicates the OMG version number of this specification document. It shall be interpreted without respect to case, and any spaces (for example, in “1.0 Beta 1”) shall be replaced with underscores.

In the event that such `pragma` declarations are nested within one another, the innermost version number specified, if any, shall be in effect. If version numbers are used with “end” declarations, those version numbers should be the same as those used with the matching “begin” declarations.

In the event that such a `pragma` “begin” declaration is not matched with a subsequent closing “end” declaration, the “begin” declaration shall be considered to continue until the end of the IDL input.

For example:

```
#pragma dds_xtopics begin 1.0_Beta_1

struct Base {
    @key long id;
};
```

```
#pragma dds_xtopics begin 1.1

struct Sub : Base {
    long another_member;
};

#pragma dds_xtopics end 1.1
#pragma dds_xtopics end 1.0_Beta_1
```

The above declarations are informative only. The behavior of an IDL compiler upon encountering them is unspecified but may include:

- Silently ignoring them.
- Issuing a warning, perhaps because it does not recognize them, or because it recognizes the pragmas but not the indicated version number.
- Halting with an error, perhaps because it recognized the pragmas and knows that it is not compliant with this specification, or because it detected a version mismatch between matching “begin” and “end” declarations.

7.3.1.2 Annotation Language

This specification makes use of different standard annotation groups defined in [IDL41]. It also proposes an alternative annotation syntax for pre-existing IDL compilers.

7.3.1.2.1 Built-in Annotations

This specification uses the following IDL annotations to model certain properties of the type system model defined in Clause 7.2.2.

In IDL an annotation may be applied to any construct or sub-construct (see Sub Clause 7.4.15.2, [IDL41]). This specification restricts the applicability of annotations to constructed types, bitmask constants, enumerated type literals, and members of aggregated types.

7.3.1.2.1.1 Member IDs

All members of aggregated types have an integral member ID that uniquely identifies them within their defining type. By default, member IDs are set automatically following a progression that starts from the most-recently specified ID (using the `@id` annotation defined in Sub Clause 8.3.1.2 in [IDL41]) or an implicit value of zero for the first constant—if there is no previous specified value—adding one with each successive member.

This behavior may be altered by two additional annotations. The `@autoid` annotation (defined in Sub Clause 8.3.1.2 in [IDL41]), which if set to `HASH` indicates that all member IDs shall be computed with a hashing algorithm, regardless of the order in which they are declared. And the `@hashid` member annotation, which provides the value to hash to generate the member ID; its definition is as follows:

```
@annotation hashid {
    string value default "";
};
```

The @hashid annotation is useful when one type is using the @autoid annotation and a new version of the type changes a member's name. The value for this annotation can be set to the old member's name, resulting in both members getting assigned the same hash value for their IDs.

If the annotation is used without any parameter or with the empty string as a value, then the Member ID shall be the hash of the member name.

7.3.1.2.1.2 Optional Members

By default, a member declared in IDL is not optional. To declare a member optional, users shall apply the @optional annotation, which is defined in Sub Clause 8.3.1.3 of [IDL41].

It is an error to declare the same member as both optional and as a key.

7.3.1.2.1.3 Key Members

By default, members declared in IDL are not considered part of their containing type's key. To declare a member as part of the key, users shall apply the @key annotation defined in Sub Clause 8.3.2.1 of [IDL41].

It is an error to declare the same member as both optional and as a key.

7.3.1.2.1.4 External Data

A member declared as external within an aggregated type indicates that it is desirable for the implementation to store the member in storage external to the enclosing aggregated type object. A suitable implementation in common programming languages may be a pointer to the member. Unless also annotated as Optional, external members shall always be present and therefore the pointer (if that is the representation used) to non-optional external members cannot be NULL. Non-optional external members can be annotated as Key.

The purpose of external data (annotated as @external) is not to facilitate graph modeling or graph (de-) serialization. If a conforming implementation encounters a graph (case #2 and #3 below), it is not required to maintain the graph structure through serialization/deserialization.

Non-normative note: Three main cases arise when using external data (1) tree structure—it is (de-) serializable (2) Diamond case—it is serializable but the bottom-most shared object may be serialized twice turning the graph into a tree. The diamond case is expected to work with some overhead. (3) Cycles—it is not serializable. However a conforming implementation is not required to warn or detect such cases.

To declare a member of an aggregated type external, apply the built-in “external” annotation to that member like this:

```
@external long my_aggregation_member;

or:

long my_aggregation_member; //@external
```

To declare the elements of a collection type external, apply the annotation to the collection declaration like this:

Sequences:

```
sequence<@external Foo, 42> sequence_of_foo;
```

Arrays:

```
Foo array_of_foo @external [42];
```

Maps:

```
map<string, @external Foo, 42> map_of_string_to_foo;
```

7.3.1.2.1.5 Enumerated Literal Values

Prior to this specification, it was impossible to indicate that objects of enumerated types could be stored using an integer size other than 32 bits. This specification uses the `@bit_bound` annotation defined in Sub Clause 8.3.4.1 of [IDL41] for this purpose.

It is important to note that the value member of the annotation may take any value from 1 to 32, inclusive, when this annotation is applied to an enumerated type.

Furthermore, prior to this specification, it was impossible to provide an explicit value for an enumerated literal. The value was always inferred based on the definition order of the literals. That behavior is still supported. However, additionally, this specification allows enumerated literals to be given explicit custom values, just as they can be in the C and C++ programming languages. This can be done by means of the `@value` annotation defined in Sub Clause 8.3.1.5 of [IDL41], which may be applied to individual literals.

It is permitted for some literals in an enumerated type to bear the `@value` annotation while others do not. In such cases, as in C and C++ enumerations, implicit values are assigned in a progression starting from the most-recently specified value (or an implicit value of zero for the first literal, if there is no previous specified value) and adding one with each successive literal.

7.3.1.2.1.6 Bitmask Positions

By default, the size of a bit mask is 32-bit. This behavior may be amended with the use of the `@bit_bound` annotation, which may set the size of the whole bit mask to a value lower or equal to 64 as specified in Sub Clause 7.4.13.4.3.3 of [IDL41].

Likewise, a bit value may be set explicitly by means of the `@position` annotation, which is defined in Sub Clause 8.3.1.4 of [IDL41].

7.3.1.2.1.7 Nested Types

By default, aggregated types and aliases to aggregated types defined in IDL are not considered to be nested types. This designation may be changed by applying the IDL `@nested` annotation to a type definition. The `@nested` annotation is defined in Sub Clause 8.3.4.3 of [IDL41].

7.3.1.2.1.8 Type Extensibility and Mutability

The extensibility kind of a type may be defined by means of a `@extensibility` annotation defined in Sub Clause 8.3.1.6 of [IDL41].

This annotation may be applied to the definitions of aggregated types. It shall be considered an error for it to be applied to the same type multiple times.

In the event that the representation of a given type does not indicate the type's extensibility kind, the type shall be considered appendable. Implementations may provide a mechanism to override this default behavior; for example, IDL compilers may provide configuration options to allow users to specify whether types of unspecified extensibility are to be considered final, appendable, or mutable.

IDL compilers shall also implement the shortcut annotations for the different extensibility kinds. That is, `@final` and `@mutable`, which defined in Sub Clauses 8.3.1.7 and 8.3.1.8 of [IDL41], as well as `@appendable`, which shall be defined as follows:

```
@annotation appendable {};
```

7.3.1.2.1.9 Must Understand Members

By default, the assignment from an object of type T2 into an object of type T1 where T1 and T2 are non-final types will ignore any members in T2 that are not present in T1. This behavior may be changed by applying the `@must_understand` annotation to a member within a type definition. The `@must_understand` annotation is defined in Sub Clause 8.3.2.2 of [IDL41].

If the `@must_understand` annotation is set to true in particular member M2 of a type T2, then the assignment to an object of type T1 shall fail if the type T1 does not define such a member.

7.3.1.2.1.10 Default Literal for Enumeration

Normally the default value for an object of a type is pre-defined based on the generic rules based on the characteristics of the type. For example, for an integer it would be the value zero and for an enumeration it is the literal with the lowest member ID.

This generic rule is not desirable in some situations. The annotation `@default_literal` allows this behavior to be changed.

```
@annotation default_literal {};
```

The application to enumerated types is illustrated in the example below:

```
enum MyEnum {  
    ENUM1,  
    ENUM2,  
    @default_literal ENUM3,  
    ENUM4  
};
```

7.3.1.2.1.11 TryConstruct Elements and Members

The construction of an object of a collection or aggregated type operates recursively; it requires constructing objects of the nested element/member types. Therefore failure to construct any object of the nested element/member type failure may impact the ability to construct the whole collection/aggregated type:

- In some cases the consequence will be that there is no object of the collection/aggregated type that can be constructed.
- In other cases the failure in the nested element/member will be mitigated and the collection/aggregated object successfully created.

The specific behavior depends on the TryConstruct behavior associated with the element or member of the type being constructed as described in 7.2.2.7.

The `@try_construct` annotation is used to explicitly set the TryConstruct behavior of element of a collection type and/or member of an aggregated type.

The IDL definition of the `@try_construct` annotation is:

```
enum TryConstructFailAction {
    DISCARD,
    USE_DEFAULT,
    TRIM
};

@annotation try_construct {
    TryConstructFailAction value default USE_DEFAULT;
};
```

As specified in 7.2.2.7 the default behavior is `DISCARD`. Therefore if the `@try_construct` annotation is not used it is the same as if it had been explicitly set to `DISCARD`. For example:

```
struct T1 {
    long important_member;
    @try_construct(DISCARD) string<4> m1;
};
```

Is the same as:

```
struct T1 {
    long important_member;
    string<4> m1;
};
```

If the annotation is specified without a value, or if the value is set to `USE_DEFAULT`, then the behavior is set to `DEFAULT` as specified in 7.2.2.7. This means the element or member will be constructed to have its default value (according to its type as described in Table 9) and does not cause the aggregated container to fail the construction.

As specified in 7.2.2.7, the TryConstruct annotation may be used in structure and union members, the union discriminator, the elements of arrays and sequences, and the key and/or values of map types.

7.3.1.2.1.11.1 TryConstruct Example 1

Assume T1 is defined:

```
struct T1 {  
    long a_long;  
    @try_construct(USE_DEFAULT) string<5> member;  
};
```

Or alternatively T1 is defined:

```
struct T1 {  
    long a_long;  
    @try_construct string<5> member;  
};
```

Assume further that T2 is defined as:

```
struct T2 {  
    long a_long;  
    string<32> member;  
};
```

In this situation if O2 is an object of type T2, and the value of the nested member object O2.member is the string “Hello World!”, then O2.member cannot construct any object of type String4 (string<5>). However since the TryConstruct behavior associated with the T1 member “member” is USE_DEFAULT, then the failure is mitigated and an O1 object of type T1 can be successfully constructed. The constructed object would have O1.member set to the empty string.

7.3.1.2.1.11.2 TryConstruct Example 2

Assume T1 and T2 are defined as:

```
struct T1 {  
    long a_long;  
    @try_construct(TRIM) string<5> member;  
};
```

```
struct T2 {  
    long a_long;  
    string<32> member;  
};
```

In this situation if O2 is an object of type T2, and the value of the nested member object O2.member is the string “Hello World!”, then the object O2.member cannot construct any object of the type of the corresponding member of T1 (string<5>). However, since the TryConstruct behavior associated with the member is TRIM, then the failure is mitigated and an object O1 of

type T1 can be successfully constructed. The constructed object would have O1.member contain the characters of O2.member that can fit on its string<5> type, that is, the string “Hello”.

7.3.1.2.1.11.3 TryConstruct Example 3

Assume T1 and T2 are defined as:

```
struct T1 {
    long a_long;
    @try_construct(TRIM) sequence<long, 4> member;
};
```

```
struct T2 {
    long a_long;
    sequence<long, 32> member;
};
```

In this situation if O2 is an object of type T2, and the value of the nested member object O2.member is the sequence of longs [1, 2, 3, 4, 5, 6, 7, 8], then the object O2.member cannot construct any object of the type of the corresponding member of T1 (sequence<long, 4>). However since the TryConstruct behavior associated with the member is TRIM, then the failure is mitigated and an object O1 of type T1 can be successfully constructed. The constructed object would have O1.member as a sequence of 4 longs containing the first four elements of O2.member.

7.3.1.2.1.11.4 TryConstruct Example 4

Assume T1 and T2 are defined as:

```
typedef string<5> String5;
struct T1 {
    long a_long;
    sequence<@try_construct(TRIM) String5, 10> member;
};
```

```
typedef string<16> String16;
struct T2 {
    long a_long;
    sequence<String16, 10> member;
};
```

In this situation if O2 is an object of type T2, and the value of the nested member object O2.member is a sequence of String16 where the first element (O2.member[0]) is “Hello World”, then the object O2.member [0] cannot construct any object of the type of the corresponding element of T1 (String5). However since the TryConstruct behavior associated with the element of the sequence is TRIM, then the failure is mitigated and an object O1 of type T1 can be

successfully constructed. The constructed object would have `O1.member[0]` as the string “Hello” (i.e. the result of trimming “Hello World!” to the length that can fit into the `String5` element type).

7.3.1.2.1.11.5 TryConstruct Example 5

Assume T1 and T2 are defined as:

```
enum T1Enum {
    ENUM1,
    @default_literal ENUM2
};

union T1 switch ( T1Enum ) {
case ENUM1:
    long e1_value;
case ENUM2:
    long e2_value;
};

enum T2Enum {
    ENUM1,
    @default_literal ENUM2,
    ENUM3
};

union T2 switch ( T2Enum ) {
case ENUM1:
    long e1_value;
case ENUM2:
    long e2_value;
case ENUM3:
    long e3_value;
};
```

In this situation if `O2` is an object of type `T2`, and the value of the discriminator is `ENUM3`, then `O2.discriminator` cannot construct an object of type `T1Enum` and as a consequence `O2` cannot construct any object of type `T1`.

However if `T1` and `T2` had been defined to have `USE_DEFAULT TryConstruct` behavior for the discriminator as in:

```
union T1 switch ( @try_construct T1Enum ) {
case ENUM1:
```

```

        long e1_value;
    case ENUM2:
        long e2_value;
};

union T2 switch (T2Enum ) {
    case ENUM1:
        long e1_value;
    case ENUM2:
        long e2_value;
    case ENUM3:
        long e3_value;
};

```

Then in this situation the failure to construct a T1Enum from O2.discriminator would be mitigated and O1.discriminator would be set to its default value (ENUM2) and O1.e1_value would be constructed from O2.e3_value. This would allow the successful construction of an O1 object of type T1.

7.3.1.2.1.12 Verbatim Text

Verbatim Text objects associated with a constructed type declaration shall be indicated using the following `@verbatim` annotation defined in Sub Clause 8.3.5.1 of [IDL41].

7.3.1.2.1.13 Non-serialized Members

By default, all members declared in IDL are serialized. To declare that a member should be omitted from serialization, apply the `@non_serialized` annotation. The equivalent definition of this type follows:

```

@annotation non_serialized {
    boolean value default TRUE;
};

```

It is an error to declare the same member as both `non_serialized` and as a key.

7.3.1.2.2 Using Built-in Annotations

The application of the annotations listed above is restricted to the elements of specified in Table 21.

Table 21 – IDL Built-in Annotations Usage

<i>Annotation</i>	<i>Applicable</i>
@id, @optional, @must_understand, @non_serialized	Structure Members
@external, @try_construct	Structure Members, Union members (except

	union discriminator)
@key	Structure Members, Union discriminator
@bit_bound	Enumerated Types, Bit Mask Types
@extensibility, @mutable, @appendable, @final, @nested	Type declarations
@default_literal, @value	Enumerated Literals
@position	Bitmask Values
@autoid	Module declarations, Structure declarations, Union declarations
@verbatim	All elements

7.3.1.2.3 Alternative Annotation Syntax

It is anticipated that it will take vendors some amount of time to implement the syntax defined in [IDL41]. During this time, existing customers may have the need to share IDL files between products that do support this specification and those that do not. In such a case, the extended annotation syntax defined here could be problematic. Therefore, this specification defines an alternative syntax for annotations that will not cause problems for pre-existing IDL compilers.

This alternative syntax uses special comments containing at-signs ('@'), much like the way Javadoc used "at" comments to attach metadata to declarations prior to the introduction of an annotation to the Java language. (For example, the conventional way to deprecate a method prior Java 5 was to place @deprecated in the documentation. In Java 5 and above, the preferred way is to use @deprecated in the source code itself, but the Javadoc-based mechanism is still supported.)

As an alternative to prefixing a declaration with an annotation, it is legal to follow the declaration with a single-line comment containing the annotation string. To distinguish such comments from regular comments, there must be no space in between the double slash ("//") and the at-sign ('@'). For example:

```
struct Gadget {
    long my_integer; //@my_member_annotation("Hello")
}; //@my_type_annotation
```

If multiple annotations are to be applied to the same element, the at-sign of each shall be preceded by a double slash and no white space. For example:

```
struct Gadget {
    long my_integer; //@my_annotation1(greeting="Hello")
    //@my_annotation2
}; //@my_type_annotation
```

7.3.1.2.4 Defining Annotations

Annotation types shall be represented as described in this clause. An annotation type is defined using the new token `@annotation`, as in the following example:

```
@annotation MyAnnotation {
    // ...
};
```

Annotation identifiers are orthogonal to any other kind of type and therefore do not conflict with other types that may use the same identifier name even when defined in the same module. This is because the application of an annotation prefixes the annotation identifier with the “@” character, see Sub Clause 7.3.1.2.5.

Recall from the Type System Model that annotation types are a form of aggregated type similar to a structure. The members of these types shall be represented using IDL members, as shown in the following example:

```
@annotation MyAnnotation {
    long my_annotation_member_1;
    double my_annotation_member_2;
};
```

Annotation members have additional constraints that are described above in the Type System Model.

Table 22 – Syntax for declaring an annotation type

@annotation <code><ann_identifier> “{”</code> <code><ann_members></code> <code>“};”</code>	Declares an annotation type containing the members <code><ann_members></code> .
struct <code><ann_identifier></code> <code>“{”</code> <code><ann_members></code> <code>“};” // @annotation</code>	The “struct” <code><ann_identifier></code> is actually an annotation type containing the members <code><ann_members></code> . The Alternative annotation syntax has been defined for backward compatibility with legacy IDL compilers.

Annotation members can take default values; these are expressed by using the keyword “default” in between the attribute name and the semicolon, followed by the default value. This value must be a valid IDL literal that is type compatible with the type of the member.

Table 23 – Syntax for members of annotation types

<code>[<pre_annotations>] <member_type></code> <code><member_name> [default</code> <code><member_value>];</code> <code>[<post_annotations>]</code>	The enclosing annotation has a member <code><member_name></code> of type <code><member_type></code> . That member may have other annotations applied to it, either before or (equivalently) after.
--	--

Consider the following example⁴. The `RequestForEnhancement` annotation indicates that a given feature should be implemented in a hypothetical system, and it provides some additional information about the requested enhancement.

```
@annotation RequestForEnhancement {
    long id;           // identify the RFE
    string synopsis;  // describe the RFE
    string engineer   default "[unassigned]"; // engineer to implement
    string date       default "[unimplemented]"; // date to implement
};
```

The specified default value may be any legal IDL literal compatible with the declared member type.

7.3.1.2.5 Applying Annotations

Annotations may be applied to any type definition or type member definition. The syntax for doing so is to prefix the definition with an at-sign ('@') and the name of the desired annotation interface. For example:

```
struct Delorean {
    Wheel wheels[4];
    float miles_per_gallon;
    @RequestForEnhancement boolean can_travel_through_time;
};
```

More than one annotation may be applied to the same element, and multiple instances of the same annotation may be applied to the same element.

Table 24 – Syntax for applying annotations

<pre>{ "@<annotation_type_name> ["("<arguments> ")"] }*</pre>	<p>Apply an annotation to a type or type member by prefixing it with an at sign ('@') and the name of the annotation type to apply. To specify the values of any members of the annotation type, include them in <i>name=value</i> syntax between parentheses.</p>
<pre>{ "//@<annotation_type_name> ["("<arguments> ")"] }*</pre>	<p>Alternately and equivalently, apply an annotation to a type or type member by suffixing it with an annotation type name using slash-slash-at ("//@") instead of the at sign by itself.</p>

Annotations can be applied to the implicit discriminator member of a union type by applying them to the discriminator type declaration in the header of the union type's definition:

```
union MyUnion switch (@MyAnnotation long) {
```

⁴ The example annotation type shown is based on one used in the Java annotation tutorial from Sun Microsystems: <http://java.sun.com/j2se/1.5.0/docs/guide/language/annotations.html>.

```

case 0:
    string member_0;
default:
    long default_member;
};

```

As with any IDL identifier, the name of an annotation and its members are *not* case-sensitive. To specify multiple annotations, place them one after another, separated by white space.

To specify values for any or all of the annotation type’s members, follow the name of the annotation with a parenthesis, and place the member values in a comma-delimited list in between them, where each list item is of the form “*member_name = member_value*.” Each value must be a compile-time constant. For example:

```

struct Delorean {
    @RequestForEnhancement(
        id = 10,
        synopsis = "Enable time travel",
        engineer = "Mr. Peabody",
        date = "4/1/3007"
    )
    boolean can_travel_through_time;
};

```

An annotation with an empty list of member values is equivalent to a member list that is omitted altogether.

Any member of the annotation interface may be omitted when the annotation is applied. If a value for a given member is omitted, and that member has a defined default value, it will take that value. If an omitted member does not have a specified default, it will take the default value specified for its type in Clause 7.2.2.4.4.5.

If an annotation interface has only a single member, the type designer is recommended to name that member “value.” In such a case, the member name may be omitted when applying the annotation. For example:

```

@annotation Widget {
    long value;
};

@Widget(5)
struct Gadget {
    // ...
};

```

7.3.1.2.6 Alternative Syntax

It is anticipated that it will take vendors some amount of time to implement this specification. During this time, existing customers may have the need to share IDL files between products that do support this specification and those that do not. In such a case, the extended annotation syntax defined here could be problematic. Therefore, this specification defines an alternative syntax for annotations that will not cause problems for pre-existing IDL compilers.

This alternative syntax uses special comments containing at-signs ('@'), much like the way Javadoc used "at" comments to attach meta-data to declarations prior to the introduction of an annotation to the Java language. (For example, the conventional way to deprecate a method prior to Java 5 was to place @deprecated in the documentation. In Java 5 and above, the preferred way is to use @Deprecated in the source code itself, but the Javadoc-based mechanism is still supported.)

As an alternative to prefixing a declaration with an annotation, it is legal to *follow* the declaration with a single-line comment containing the annotation string. To distinguish such comments from regular comments, there must be no space in between the double slash ("//") and the at-sign ('@'). For example:

```
struct Gadget {  
    long my_integer;    //@MyMemberAnnotation("Hello")  
}; //@MyTypeAnnotation
```

If multiple annotations are to be applied to the same element, the at-sign of each shall be preceded by a double slash and no white space. For example:

```
struct Gadget {  
    long my_integer; //@MyAnnotation1(greeting="Hello")  
                    //@MyAnnotation2  
}; //@MyTypeAnnotation
```

7.3.1.3 Constants and Expressions

IDL allows the declaration of global and namespace-level constant values. It also allows the use of compile-time mathematical expressions, which may include constants, enumeration values, and numeric literals. Such declarations and expressions remain legal IDL. However, they are not reflected directly in the Type System specified here, which assumes that all compile-time-constant values have already been evaluated.

7.3.1.4 Primitive Types

The primitive types specified here directly correlate to the primitive types that already exist in IDL.

Table 25 – IDL primitive type mapping

<i>Type System Model Type</i>	<i>IDL Type</i>	<i>Type System Model Type</i>	<i>IDL Type</i>
Int16	short	Float64	double

UInt16	unsigned short	Float128	long double
Int32	long	Char8	char
UInt32	unsigned long	Char16	wchar
Int64	long long	Boolean	boolean
UInt64	unsigned long long	Byte	octet
Float32	float		

7.3.1.5 Alias Types

Aliases as described in this specification are fully compatible with the IDL `typedef` construct.

7.3.1.6 Array and Sequence Types

Arrays and sequences as described in this specification are fully compatible with the IDL constructs of the same names.

7.3.1.7 String Types

The string container defined by this specification has two element types for which the behavior is defined: `Char8` and `Char16`. Strings of `Char8` shall be represented by the IDL type `string`. Strings of `Char16` shall be represented by the IDL type `wstring`. In either case, any bound shall be retained.

7.3.1.8 Enumerated Types

Enumerations and bitmasks as described in this specification are fully compatible with the IDL constructs of the same name.

7.3.1.9 Map Types

Map types as described in this specification are fully compatible with the IDL constructs of the same name defined in the Extended Data-Types Building Block of [IDL41].

Structures as defined by this specification are fully compatible with the IDL constructs of the same name.

7.3.1.10 Structure Types

Structures as described in this specification are in this specification are fully compatible with the IDL constructs of the same name.

7.3.1.11 Union Types

Unions as described in this specification are fully compatible with the IDL constructs of the same name. Compliant IDL parsers shall implement the Building Block Extended Data-Types of [IDL41], which adds support for `Byte` (`octet`) and `Char16` (`wchar`) type discriminators.

7.3.2 XML Type Representation

Types may be defined in an easy-to-read, easy-to-process XML format. This format is defined by an XML schema document (XSD) and a set of semantic rules, which are discussed below.

The XML namespace of the XML Type Representation shall be `http://www.omg.org/dds`.

Design Rationale (non-normative)

The XML Type Representation very much resembles a translation of the grammar of the IDL Type Representation directly into XML. The largest change from such a straightforward translation is that the “built-in annotations” from the IDL Type Representation are here represented as first-class XML constructs—a luxury that is feasible here because this Representation does not predate the definition of the corresponding modeling concepts.

7.3.2.1 Type Representation Management

This Type Representation provides several features that do not directly impact or reflect the Type System. However, they provide capabilities that are necessary or convenient for the organization and management of type declarations. These features are described in this clause.

7.3.2.1.1 File Inclusion

As in IDL, files may include other files. Such inclusions shall not be considered semantically meaningful with respect to the Type System Model, but they can be useful as a code maintenance tool.

A file inclusion specified as in this Type Representation shall be considered equivalent to an IDL `#include` of the same file. A formal definition is in “Annex A: XML Type Representation Schema.” The following is a non-normative example:

```
<dds:types xmlns:dds="http://www.omg.org/dds">
  <dds:include file="my_other_types.xml"/>
</dds:types>
```

Conformant Type Representation compilers need not support the inclusion of files of other Type Representations from within an XML Type Representation document. For example, conformant Type Representation compilers need not support the inclusion of IDL files from XML files.

Design Rationale (non-normative)

XML provides other mechanisms to include one file within another—for example, by defining custom entities. However, these mechanisms cannot provide functionality equivalent to the `#include` of IDL because of when they are interpreted during the XML parsing process.

For example, suppose a type `X` defined in `X.xml` and a type `Y` defined in `Y.xml` both depend on a type `Z` defined in `Z.xml`. Suppose further that an application wishes to use these three types using their Plain Language Bindings in the C programming language. If `X.xml` and `Y.xml` include `Z.xml` using an XML entity definition, this definition will be expanded by the XML parser (upon which the code generator is presumably implemented), and the code generator will never know of the existence of `Z.xml`. It will instead encounter two definitions of `Z`, and the application will fail to build because of multiply defined symbols.

As an alternative, the mechanism described here allows the code generator to *observe* the intention to include `Z.xml` and generate `#include <Z.h>`, avoiding the multiple definition problem.

7.3.2.1.2 Forward Declarations

As in IDL, C, and C++, a usage of a type must be preceded by a declaration of that type. Therefore, as those languages do, this Type Representation provides for *forward declarations* of types. These declarations are provided for the convenience of code generator implementations; they shall have no representation in the Type Representation Model.

A forward declaration as described in this Type Representation shall be considered semantically equivalent to an IDL forward declaration. A formal definition is in “Annex A: XML Type Representation Schema.” The following is a non-normative example:

```
<dds:types xmlns:dds="http://www.omg.org/dds">
  <dds:forward_decl kind="struct" name="MyStructure"/>
</dds:types>
```

7.3.2.1.3 Constants

As in the IDL Type Representation, the XML Type Representation supports declaration of compile-time constant values. Specifically, the string specified in the `value` attribute described below shall have the same syntax as the **<const_exp>** production in the IDL grammar [IDL41].

Constants can appear at the top level of a Type Representation file, within a module, or—as in an IDL `valuetype`—within a structure declaration.

Constants are not reflected directly in the Type System. Instead, mathematical expressions shall be considered to be evaluated at compile time.

The following is a non-normative example:

```
<dds:types
xmlns:dds="http://www.omg.org/ptc/2011/01/07/XML_Type_Representation">
  <dds:const name="MY_CONSTANT" type="int32" value="2 + 3"/>
</dds:types>
```

7.3.2.2 Basic Types

This Type Representation represents type names with a combination of XML attributes, defined according to the following pattern:

- A “type” attribute, typed by an enumeration `allTypeKind`, indicates whether the type is “basic” (i.e., is a primitive or string)—and if so, which one—or if it is “non-basic” (i.e., any other type).

Design rationale: As even basic types have identifier names, the use of the `allTypeKind` enumeration does not add to the expressiveness of this Type Representation. However, since primitive types are used frequently, the enumeration allows XML editors to provide context-sensitive completions, improving the user experience.

- A “non-basic type name” attribute indicates the name of the type if it is a non-basic type. It is an error to include this attribute if the `type` attribute does not indicate a non-basic type.
- If the type is a collection type, additional attributes describe its bound(s); see below.

The names of the basic types in this Type Representation have been chosen to resemble terse versions of the corresponding names in the Type System Model.

Table 26 – Primitive and string type names in the XML Type Representation

<i>Type System Model Name</i>	<i>XML Type Representation Name</i>
Boolean	boolean
Byte	byte
Char8	char8
Char16	char16
Int32	int32
UInt32	uint32
Int16	int16
UInt16	uint16
Int64	int64
UInt64	uint64
Float32	float32
Float64	float64
Float128	float128
String<Char8, ...>	string
String<Char16, ...>	wstring

7.3.2.3 String Types

As described above, strings (whether of narrow or wide characters) are considered to be basic types in this Type Representation. Nevertheless, the description of their bounds requires additional attributes.

The `stringMaxLength` attribute, if present, indicates the string’s bound. If the attribute is omitted, the string shall be considered unbounded.

The presence of this attribute is legal only when a member’s type is a string, a wide string, or an alias to string or wide string. The following examples are non-normative:

```
<struct name="MyStructure">
  <member name="unbounded_string_1" type="string"/>
</struct>
```

```

    <member name="unbounded_string_2" type="string" stringMaxLength="-1"/>
    <member name="bounded_string" type="string"
        stringMaxLength="2 + MY_CONSTANT"/>
</struct>

```

7.3.2.4 Collection Types

The element type identified by the `type` and `nonBasicTypeName` attributes correspond to the type of a member itself when the member identifies a single value, to the element type when the member is of a sequence or array collection, or to the “value” type of map collection if the member is of a map type. This clause and its sub clauses summarize these rules; the formal grammar can be found in “Annex A: XML Type Representation Schema.”

Collection bounds are indicated by attributes named according to the convention

`<collection>MaxLength`: `stringMaxLength`, `sequenceMaxLength`, and `mapMaxLength`. The types of these attributes are strings, not integers: the values of these attributes may be any constant expression as defined by the **<const_exp>** production in the IDL grammar [IDL41]. The literal expression “-1” shall indicate an unbounded collection; no other “negative” value is permitted.

The `element_external` property of the Type System Model shall be represented by an attribute `external`.

7.3.2.4.1 Array Types

The presence of the `arrayDimensions` attribute shall indicate that given member is an array. Array dimensions are represented as a comma-delimited list of dimension bounds in the same order in which those bounds would be given in IDL. Whitespace is allowed around each bound and is not significant.

Compile-time-constant mathematical expressions are also permitted; their syntax shall be defined by the **<const_exp>** production in the IDL grammar [IDL41]. As in the IDL Type Representation, such expressions are not expressed directly in the Type System Model but are evaluated first. For example, the following are all valid:

- `arrayDimensions="1"`
- `arrayDimensions="2, MY_CONSTANT + 3"`
- `arrayDimensions=" 6,2, 3 "`

For example:

```

<struct name="MyStructure">
    <member name="my_array_of_42_integers" type="int32" arrayDimensions="42"/>
</struct>

```

7.3.2.4.2 Sequence Types

The `sequenceMaxLength` attribute, if present, shall indicate that the member is of a sequence type.

The following is a non-normative example:

```
<struct name="MyStructure">
  <member name="my_unbounded_sequence_of_integers" type="int32"
    sequenceMaxLength="-1"/>
  <member name="my_bounded_sequence_of_structures" type="nonBasic"
    nonBasicTypeName="MyOtherStructure"
    sequenceMaxLength="6 * 3"/>
</struct>
```

7.3.2.4.3 Map Types

Map types must include the following additional information:

- The map’s bound, if any, shall be indicated by the `mapMaxLength` attribute. This attribute is required for all map types.
- The type of the map’s “key” elements shall be indicated by the `mapKeyType` attribute. This attribute is required for all map types. This attribute is exactly parallel to the `type` attribute (which describes the type of the map’s “value” elements): it indicates whether the “key” elements of the map are of a basic or non-basic type and, if basic, which basic type. If the type is non-basic, the `mapKeyNonBasicTypeName` attribute is also required and is parallel to the `nonBasicTypeName` attribute. If the “key” type is basic, the `mapKeyNonBasicTypeName` attribute is not allowed.
- Only if the map’s “key” type is a string type, the attribute `mapKeyStringMaxLength`, if present, shall indicate the bound of that string type. If the “key” type is a string type, and this attribute is omitted, the string shall be considered unbounded. If the “key” type is not a string type, this attribute is not allowed.

The following is a non-normative example:

```
<struct name="MyStructure">
  <member name="my_unbounded_maps_of_integers_to_floats" type="int32"
    mapKeyType="float32"
    mapMaxLength="-1"/>
  <member name="my_bounded_map_of_strings_to_structures"
    mapKeyType="string"
    mapKeyStringMaxLength="128"
    type="nonBasic"
    nonBasicTypeName="MyOtherStructure"
    mapMaxLength="6 * 3"/>
</struct>
```

7.3.2.4.4 Combinations of Collection Types

A type may be a sequence of arrays, a map of strings to sequences, or some other complex combination of collection types. It's therefore important to understand, if some *combination* of `sequenceMaxLength` and `mapMaxLength` are present, which takes precedence. The following list is ordered from most-tightly-binding to least-tightly-binding:

- Sequence designations, including `sequenceMaxLength`
- Array designations, including `arrayDimensions`
- Map designations, including `mapMaxLength`.

To indicate a type composed in a different order (for example, a sequence of arrays), it is necessary to interpose an alias definition.

For example, a member specifying all of these would define a map whose values are arrays of sequences of strings. Further examples follow:

```
<struct name="MyStructure">
  <member name="my_array_of_strings"
    type="string"
    stringMaxLength="-1"
    arrayDimensions="20"/>
  <member name="my_array_of_sequences_of_integers"
    type="int32"
    sequenceMaxLength="6 * 3"
    arrayDimensions="20"/>
</struct>
```

7.3.2.5 Aggregated Types

Aggregated types include those types that define internal named members taking per-instance values: annotations, structures, and unions.

The Type System defines a number of properties for aggregated types and their members:

- `extensibility_kind`
- `nested`
- `key`
- `optional`
- `must_understand`, etc.

The IDL Type Representation is based on IDL, which provides no syntax to provide values for these attributes; therefore, that Type Representation makes use of built-in annotations for this purpose. In contrast, the XML Type Definition is able to express these properties directly.

For example, structures and unions may indicate whether they are appendable/mutable and/or nested types:

```
<struct name="MyStructure" extensibility="mutable" nested="true">
  ...
</struct>
```

In the event that the representation of a given type does not indicate the type's extensibility kind, an implementation may make its own determination. In particular, type representation compilers shall provide configuration options to allow users to specify whether types of unspecified extensibility will be considered final, appendable, or mutable.

7.3.2.5.1 Structures

Structures contain four kinds of declarations:

- Applied annotations
- Verbatim text
- Members
- Constants

Constants and applied annotations are described above. The other elements are described in the sections below.

7.3.2.5.1.1 Verbatim Text

As described in Clause 7.2.2.4.5, types may store blocks of text to be used by Type Representation compilers. These are represented within a structure's declaration as shown in the following non-normative example:

```
<struct name="MyStructure">
  <verbatim language="Java" placement="before-declaration">
    /**
     * This is a JavaDoc comment.
     */
  </verbatim>
  ...
</struct>
```

7.3.2.5.1.2 Members

Each structure type shall include one or more members. Each member of a structure type can indicate individually whether or not it is a key member and whether or not it is an optional member.

```
<struct name="structMemberDecl">
  <member name="my_key_field" type="int32" key="true" optional="false"/>
</struct>
```

7.3.2.5.1.3 Inheritance

A structure declaration's `baseType` attribute indicates the name of the structure's base type, if any; if it is omitted, then the structure has no base type. For example:

```
<struct name="MyStructure" baseType="MyOtherStructure">
  ...
</struct>
```

7.3.2.5.2 Unions

In addition to the `annotate` and `verbatim` elements they share with other aggregated types (see above), unions contain two kinds of members: exactly one discriminator member (identified by a `discriminator` element) and one or more cases (identified by `case` members). The discriminator member must be declared before the others.

Each case of a union contains one or more discriminator values (`caseDiscriminator` elements) and one data member. A case discriminator is a string expression, the syntax of which shall be defined by the **<const_exp>** production in the IDL grammar [IDL41]. The literal "default" is also allowed; it indicates that the corresponding case is the default case—there can only be one such within a given union declaration.

For example:

```
<union name="MyUnion">
  <discriminator type="int32"/>
  <case>
    <caseDiscriminator value="1"/>
    <caseDiscriminator value="2"/>
    <member name="small_value" type="float32"/>
  </case>
  <case>
    <caseDiscriminator value="default"/>
    <member name="large_value" type="float64"/>
  </case>
</union>
```

The example above is equivalent to the following IDL type:

```
union MyUnion switch (long) {
  case 1:
  case 2:
    float small_value;
  default:
    double large_value;
};
```

7.3.2.6 Aliases

Alias definitions are defined in `typedef` elements. They have syntax very similar to that of structure members.

For example:

```
<typedef name="MyAliasToSequenceOfStructures"
  type="nonBasic"
  nonBasicTypeName="MyStructure"
  sequenceMaxLength="16"/>
```

7.3.2.7 Enumerated Types

7.3.2.7.1 Enumerations

Enumerations consist of a list of enumeration literals, each of which has a name and a value. The syntax of the value shall be defined by the **<const_exp>** production in the IDL grammar [IDL41]. If the value is omitted, it shall be assigned automatically.

For example:

```
<enum name="MyEnumeration" bitBound="16">
  <enumerator name="LITERAL_1" value="0"/>
  <enumerator name="LITERAL_2" value="0+1"/>
  <enumerator name="LITERAL_3"/>
</enum>
```

7.3.2.7.2 Bitmasks

A bitmask type defines a sequence of flags, each of which shall identify one of the bits in the bitmask.

For example:

```
<bitmask name="MyBitmask" bitBound="64">
  <flag name="FIRST_BIT" position="0"/>
  <flag name="SECOND_BIT" position="1"/>
</bitmask>
```

7.3.2.8 Modules

A module groups type declarations and serves as a namespace for those definitions.

```
<module name="MyModule1">
  <struct name="MyStructure">
    <member name="my_member" type="int64"/>
  </struct>
</module>
```

```

<module name="MyModule2">
  <struct name="MyStructure">
    <member name="my_member" type="nonBasic"
      nonBasicTypeName="MyModule1::MyStructure"/>
  </struct>
</module>

```

7.3.2.9 Annotations

There are two primary declarations pertaining to annotations: annotation types and the applications of them to types and type members, specifying values for the annotation’s own members.

The following is a non-normative example:

```

<annotation name="MyAnnotation">
  <member name="widgets" type="int32"/>
</annotation>

<struct name="MyStructure">
  <annotate name="MyAnnotation">
    <member name="widgets" value="5"/>
  </annotate>
  ...
</struct>

```

7.3.3 XSD Type Representation

Types can be defined using an XML schema document (XSD). The format is based on the standard IDL mapping to XSD [IDL-XSD]. An XSD Representation of a given type shall be as if the OMG-standard IDL mapping to XSD were applied to the IDL Representation of the type as defined in Clause 7.3.1. That mapping is augmented as follows to address IDL extensions defined by this specification. The resulting XSD representation may be embedded within a WSDL file or may occur as an independent XSD document.

XML Schema documents intended for use with DDS, like any XML Schema documents, may declare a target namespace for the elements and attributes they define. Valid documents conforming to such schemas (i.e. serialized DDS samples; see Clause 7.4.4, “XML Data Representation”) must respect such namespaces, if any.

7.3.3.1 Annotations

It is possible to both define and apply annotations using the XSD Type Representation; these tasks shall be accomplished using XSD Annotations. (To avoid confusion, for the remainder of this clause, an annotation as defined by the Type System Model in this document will be referred to as an “OMG Annotation.” An annotation as defined by the XML Schema specification shall be referred to as an “XSD Annotation.”)

7.3.3.1.1 Defining Annotation Types

OMG Annotation types shall be defined using XSD-standard `complexType` definitions. Any `complexType` definition immediately containing an XSD Annotation with an `appInfo` element having a `source` attribute value of `http://www.omg.org/Type/Annotation/Definition` shall be considered to be an OMG Annotation. Such `complexType` definitions, henceforth referred to as “Annotation `complexType` Definitions” shall conform to the structure defined in this clause.

Each attribute of an Annotation `complexType` Definition shall define a member of the corresponding OMG Annotation type:

- The `name` of the attribute shall specify the name of the OMG Annotation type member.
- The `type` of the attribute shall specify the name of the type of the OMG Annotation type member.
- A `default` value, if present, shall specify the default value of the OMG Annotation type member.

The meanings of any sub-elements defined for an Annotation `complexType` Definition are unspecified. The following example provides equivalent definitions for an OMG Annotation type in both IDL and XSD.

Table 27– XSD annotation example

<i>IDL</i>	<i>XSD</i>
<pre>@annotation my_annotation { long widgets; double gadgets default 42.0; };</pre>	<pre><xsd:complexType name="my_annotation"> <xsd:annotation> <xsd:appInfo source="http://www.omg.org/Type/Annotation/Definition"/> </xsd:annotation> <xsd:attribute name="widgets" type="xsd:int"/> <xsd:attribute name="gadgets" type="xsd:double" default="42.0"/> </xsd:complexType></pre>

7.3.3.1.2 Applying Annotations

OMG Annotations shall be applied to a definition by declaring, immediately within that definition’s XML element, an XSD Annotation containing an `appInfo` with its `source` attribute set to `http://www.omg.org/Type/Annotation/Usage`. The structure of such an `appInfo` element shall conform to that defined in this clause.

The `appInfo` element shall contain an element `annotate` for each OMG Annotation to be applied. For syntactic validation purposes, the definition of the `annotate` element shall be as follows:

```

<xsd:schema targetNamespace="http://www.omg.org/Type"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xsd:complexType name="annotate">
    <xsd:attribute name="type" type="xs:string" use="required"/>
    <xsd:anyAttribute processContents="skip"/>
  </xsd:complexType>
  ...
</xsd:schema>

```

However, for semantic validation purposes, the `annotate` element shall contain attribute values corresponding to any subset of the attributes defined by the OMG Annotation type indicated by its required `type` attribute.

In the following example, the OMG Annotation `MyAnnotation` defined in the previous example is applied to a structure definition:

```

<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:omg="http://www.omg.org/Type"
  xmlns:tns="http://www.omg.org/IDL-Mapped/"
  targetNamespace="http://www.omg.org/IDL-Mapped/">
  <xsd:complexType name="MyStructure">
    <xsd:annotation>
      <xsd:appInfo source="http://www.omg.org/Type/Annotation/Usage">
        <omg:annotate omg:type="MyAnnotation" widgets="12"
          gadgets="75.0"/>
      </xsd:appInfo>
    </xsd:annotation>
  </xsd:complexType>
</xsd:schema>

```

7.3.3.1.3 Built-in Annotations

Unless otherwise noted, those Type System concepts represented with built-in annotations in the IDL Type Representation shall be represented by equivalent built-in annotations in this Type Representation.

7.3.3.2 Structures

The representations of structures and their members shall be augmented as described below.

7.3.3.2.1 Inheritance

The subtype shall extend its base type using an XSD `complexContent` element. For example, the following types in the IDL Type Representation and XSD Type Representation are equivalent:

Table 28 – XSD structure inheritance example

<u>IDL</u>	<u>XSD</u>
<pre> struct MyBaseType { long inherited_member; }; struct MyExtendedType : MyBaseType { long new_member; }; </pre>	<pre> <xs:complexType name="MyBaseType"> <xs:sequence> <xs:element name="inherited_member" type="xs:int"/> </xs:sequence> </xs:complexType> <xs:complexType name="MyExtendedType"> <xs:complexContent> <xs:extension base="MyBaseType"> <xs:sequence> <xs:element name="new_member" type="xs:int"/> </xs:sequence> </extension> </xs:complexContent> </xs:complexType> </pre>

7.3.3.2 Optional Members

Optional members of an aggregated type shall be indicated with a `minOccurs` attribute value of 0 instead of 1. For example:

```

<xsd:complexType name="MyType">
  <xsd:sequence>
    <xsd:element name="my_int" minOccurs="0" maxOccurs="1" type="xsd:int"/>
  </xsd:sequence>
</xsd:complexType>
        
```

7.3.3.3 Nested Types

For each type T that is *not* a nested type, the schema shall define an XML element of that type suitable for use as a document root. The name of this element shall be the fully qualified name of T.

For example, for the structure “MyStructure” in the module “MyModule” (named “MyModule.MyStructure” in this Type Representation) the schema shall include a declaration like the following:

```

<xs:element name="MyModule.MyStructure" type="MyModule.MyStructure"/>
        
```

7.3.3.4 Maps

A map declaration is superficially like a structure declaration; however, the XSD `sequence` declaration specifies a `maxOccurs` multiplicity equal to the bound of the map (or unbounded if

the map is unbounded). The map elements are represented by elements named `key` and `value`, each of which must occur exactly once for each iteration of the sequence.

For example, the following is a map of integers to floating-point numbers with a bound of 32:

```
<xsd:complexType name="MyMap">
  <xsd:sequence maxOccurs="32">
    <xsd:element name="key" minOccurs="1" maxOccurs="1" type="xsd:int"/>
    <xsd:element name="value" minOccurs="1" maxOccurs="1"
      type="xsd:double"/>
  </xsd:sequence>
</xsd:complexType>
```

7.3.4 Representing Types with `TypeIdentifier` and `TypeObject`

Any possible type within the XTYPES type system is uniquely identified by a `TypeIdentifier`. In the case of simple types such as primitive types, string, or certain sequences of primitive types, the `TypeIdentifier` completely describes the `Type`. For more complex types, the `TypeIdentifier` only identifies the type and its full description uses a `TypeObject`.

See “Annex B: Representing Types with `TypeObject`” for the formal definition of the `TypeIdentifier` and `TypeObject` types.

7.3.4.1 Plain Types

This specification uses the term **Plain Collection** type to refer to anonymous collection types (array, sequence, and map) that have no annotations beyond `@external` and `@try_construct`.

This specification uses the term **Plain** type to refer to primitive types and plain collection types. The remaining types are called **Non-Plain** types.

Plain types only have a `TypeIdentifier`. Non-plain types have both a `TypeIdentifier` and a `TypeObject`.

7.3.4.2 Type Identifier

The type identifier provides a unique way to identify each type within the XTYPES type system. More precisely it identifies each equivalence class of types, see Clause 7.3.4.6.

The definition of the type identifier uses the structure `TypeIdentifier` declared in IDL; see “Annex B: Representing Types with `TypeObject`”.

`TypeIdentifier` is a discriminated union allowing the format of the identifier to vary depending on the type. Table 29 below lists the `TypeIdentifier` discriminator values and their use.

Table 29 – Formats and interpretation of the `TypeIdentifier`

<i>TypeIdentifier discriminator value</i>	<i>Types</i>	<i>Notes</i>
TK_NONE	N/A	Invalid identifier

TK_BOOLEAN, TK_BYTE, TK_INT16, TK_INT32, TK_INT64, TK_UINT16, TK_UINT32, TK_UINT64, TK_FLOAT32, TK_FLOAT64, TK_FLOAT128, TK_CHAR8, TK_CHAR16	Primitive Types	Plain Type. No TypeObject Fully described by the discriminator. No further information in TypeIdentifier.
TI_STRING8_SMALL, TI_STRING8_LARGE	String Types	Plain Type. No TypeObject Fully described by the discriminator and the bound of the string. The SMALL discriminators have a bound represented as an octet. It is used for unbounded strings or strings with bounds smaller than 256. The LARGE discriminators are used for the remaining strings
TI_STRING16_SMALL, TI_STRING16_LARGE	Wide String types	Plain Type. No TypeObject Fully described by the discriminator and the bound of the string. SMALL and LARGE indicate representation of bound.
TI_PLAIN_SEQUENCE_SMALL, TI_PLAIN_SEQUENCE_LARGE	Plain sequence Collection	Plain Type. No TypeObject TypeIdentifier contains maximum length of sequence and the TypeIdentifier of element. SMALL and LARGE indicate representation of maximum length.
TI_PLAIN_ARRAY_SMALL, TI_PLAIN_ARRAY_LARGE	Plain array Collection	Plain Type. No TypeObject TypeIdentifier contains array dimensions and the TypeIdentifier of element. SMALL and LARGE indicate representation of dimensions.
TI_PLAIN_MAP_SMALL, TI_PLAIN_MAP_LARGE	Plain map Collection	Plain Type. No TypeObject TypeIdentifier contains length of map and the TypeIdentifier of key and element. SMALL and LARGE indicate

		representation of maximum length.
TI_STRONGLY_CONNECTED_COMPONENT	Types with mutual dependencies on other types	Not plain type. Has TypeObject. Uses a Hash computed on the TypeObjects of the set of mutually-dependent types. See clause 7.3.4.8.
EK_COMPLETE	Not mutually dependent on other types	Not plain type. Has TypeObject. Uses a Hash of the Complete TypeObject that describes the type. See 7.3.4.3
EK_MINIMAL	Not mutually dependent on other types	Not plain type. Has TypeObject. Uses a Hash of the Minimal TypeObject that describes the type. See 7.3.4.4.
TK_ANNOTATION	Annotation Declaration	Not plain type. Has TypeObject. Uses Hash of the TypeObject representation of the Annotation declaration
TI_EXTENDED		Reserved for future extensions

7.3.4.3 Complete TypeObject

The Complete TypeObject is a type representation with the same expressive power as the IDL (7.3.1, XML (7.3.2), and XSD (7.3.3) representations. Any non-plain type represented in IDL can be converted to the Complete TypeObject representation and back to IDL with no information loss, other than formatting (e.g. presence of whitespace).

The Complete TypeObject provides an alternative representation of types suitable for programming and tooling.

The complete TypeObject is defined by its IDL representation; see the declaration of structure CompleteTypeObject in “Annex B: Representing Types with TypeObject”.

7.3.4.4 Minimal TypeObject

The Minimal TypeObject provides a compact way to represent the type information relevant for a remote application to determine type assignability. This representation does not include information on the type that would not impact type assignability. For example user-defined annotations or the order of members for types with extensibility kind MUTABLE.

The Minimal TypeObject reduces the amount of information that applications need to send on the network in order to check type assignability between DataWriters and DataReaders.

The complete TypeObject is defined by its IDL representation; see the declaration of structure MinimalTypeObject in “Annex B: Representing Types with TypeObject”.

7.3.4.5 TypeObject serialization

The serialization of a `TypeObject` shall happen in accordance to its IDL declaration and the general serialization rules defined in this specification (see Clause 7.4) for XCDR encoding version 2. Additional restrictions are placed such that the serialized result is bitwise identical independently of the vendor or platform where the serialization occurs. Specifically:

- The serialization shall use Little Endian encoding.
- The elements in `AnnotationParameterSeq` shall be ordered in increasing values of their `paramname_hash` `_typeid`.
- The elements in `AppliedAnnotationSeq` shall be ordered in increasing values of their `annotation_typeid`.
- The elements in `CompleteStructMemberSeq` shall be ordered in increasing values of the `member_index`.
- The elements in `MinimalStructMemberSeq` shall be ordered in increasing values of the `member_index`.
- The elements in `CompleteUnionMember` shall be ordered in increasing values of the `member_index`.
- The elements in `MinimalUnionMember` shall be ordered in increasing values of the `member_index`.
- The elements in `CompleteAnnotationMemberSeq` shall be ordered in increasing values of the `member_index`.
- The elements in `MinimalAnnotationMemberSeq` shall be ordered in increasing values of the `member_name` `hash`.
- The elements in `CompleteEnumeratedLiteralSeq` shall be ordered in increasing values of their numeric value.
- The elements in `MinimalEnumeratedLiteralSeq` shall be ordered in increasing values of their numeric value.
- The elements in `CompleteBitflagSeq` shall be ordered in increasing values of their `position`.
- The elements in `MinimalBitflagSeq` shall be ordered in increasing values of their `position`.
- The elements in `CompleteBitfieldSeq` shall be ordered in increasing values of their `position`.
- The elements in `MinimalBitfieldSeq` shall be ordered in increasing values of their `position`.

7.3.4.6 Classification of TypeIdentifiers

7.3.4.6.1 Fully-descriptive TypeIdentifiers

Some `TypeIdentifiers` do not involve computing the hash of any `TypeObject`. These are called Fully-descriptive `TypeIdentifiers` because they fully describe the `Type`. These are:

- The `TypeIdentifiers` for Primitive and String types.
- The `TypeIdentifiers` of plain collections where the element (and key) `TypeIdentifier` a fully descriptive `TypeIdentifier`. They are recognized by the contained `PlainCollectionHeader` having `EquivalenceKind` set to `EK_BOTH`.

7.3.4.6.2 Hash TypeIdentifiers

Some `TypeIdentifiers` include within (directly or indirectly) hashes of one of more `TypeObjects`. These are called HASH `TypeIdentifiers`. These are:

- Those with discriminator `EK_MINIMAL`, `EK_COMPLETE`, or `TI_STRONG_COMPONENT`
- Those with discriminator `TI_PLAIN_SEQUENCE_SMALL`, `TI_PLAIN_SEQUENCE_LARGE`, `TI_PLAIN_ARRAY_SMALL`, `TI_PLAIN_ARRAY_LARGE`, `TI_PLAIN_MAP_SMALL`, or `TI_PLAIN_MAP_LARGE` where the contained `PlainCollectionHeader` has `EquivalenceKind` `EK_MINIMAL` or `EK_COMPLETE`.

In contrast to the Fully-descriptive Identifiers HASH identifiers only identify a `Type` but do not provide a complete description of the type without the auxiliary `TypeObjects` whose hashes are included in the `TypeIdentifier`.

HASH `TypeIdentifiers` are further classified along two orthogonal dimensions:

- Direct vs. Indirect. This classification looks at the nature of their dependency on the `TypeObjects`.
- Minimal vs Complete. This classification looks at the kind of `TypeObjects` involved.

7.3.4.6.3 Direct Hash TypeIdentifiers

These are the HASH `TypeIdentifiers` with discriminator `EK_MINIMAL`, `EK_COMPLETE`, or `TI_STRONG_COMPONENT`.

7.3.4.6.4 Indirect Hash TypeIdentifiers

These are the HASH for plain collections that have elements identified using a hash `TypeIdentifiers`. They are distinguished by:

1. Having discriminator `TI_PLAIN_SEQUENCE_SMALL`, `TI_PLAIN_SEQUENCE_LARGE`, `TI_PLAIN_ARRAY_SMALL`, `TI_PLAIN_ARRAY_LARGE`, `TI_PLAIN_MAP_SMALL`, or `TI_PLAIN_MAP_LARGE`.

2. Having the contained PlainCollectionHeader with EquivalenceKind EK_MINIMAL or EK_COMPLETE.

7.3.4.6.5 Minimal Hash TypeIdentifiers

These are HASH TypeIdentifiers that involve hashing serialized MINIMAL TypeObjects. They consist of:

- those with discriminator EK_MINIMAL
- those with discriminator TI_STRONG_COMPONENT where the contained TypeObjectId has discriminator EK_MINIMAL.
- those for plain collections where the contained PlainCollectionHeader EquivalenceKind is EK_MINIMAL

7.3.4.6.6 Complete Hash TypeIdentifiers

These are HASH TypeIdentifiers that involve hashing serialized COMPLETE TypeObjects. They consist of:

- those with discriminator EK_COMPLETE
- those with discriminator TI_STRONG_COMPONENT where the contained TypeObjectId has discriminator EK_COMPLETE.
- those for plain collections where the contained PlainCollectionHeader EquivalenceKind is EK_COMPLETE

7.3.4.7 Type Equivalence

A distributed type system where types can be defined at different locations using different representations leads to the need of defining equivalence relations between types.

In set theory an “equivalence” relation is one satisfying the reflexive, symmetric, and transitive properties. Using the “~” sign to represent the relation, the three properties can be expressed as:

- Reflexive: $T \sim T$ for every type “T” in the set of possible types.
- Symmetric: $T_1 \sim T_2$ implies $T_2 \sim T_1$
- Transitive: $T_1 \sim T_2$ and $T_2 \sim T_3$ implies $T_1 \sim T_3$

An equivalence relation partitions a set into disjoint subsets (equivalence classes) where each contains all the elements that are equivalent to each other. Being a “partition” each element belongs to exactly one of the equivalence classes.

An equivalence relation between types captures the intuitive notion that the related types “behave the same way” under a certain set of operations or use cases> because of this they can be considered to be “the same” from the perspective of those operations/use-cases.

When defining two equivalence relations R1 and R2 on the same set it may be the case that all elements that are equivalent under (R1) are also equivalent under the other (R2). In this case it is said that R1 is finer than R2, or alternatively that R2 is coarser than R1.

When this happens the finer relationship (R1) further partitions each equivalence class of the coarser (R2) in its own finer R1-equivalence classes. Said differently elements considered equivalent according to R2 may be differentiated by the R1 relation.

This specification defines two equivalence relations between types: **Complete** and **Minimal**.

- **Complete equivalence** relates types that can be considered the same for all practical uses of the type system, including code generation or displaying type information to the user.
- **Minimal equivalence** relates types that can be considered the same with regards to the type compatibility/assignability between a DataWriter and a DataReader as well as with regards to the data objects published by the DataWriter and received by the DataReader.

The formal definition of these equivalence relations is done in terms of TypeIdentifiers and TypeObjects.

- Two types are equivalent according to the Complete equivalence relation if and only if either they have equal Fully-Descriptive TypeIdentifiers, or else they have equal Complete TypeIdentifiers.
- Two types are equivalent according to the Minimal equivalence relation if and only if either they have equal Fully-Descriptive TypeIdentifiers, or else they have equal Minimal TypeIdentifiers.

From the definition of the Complete and Minimal TypeIdentifier it is clear that two types that are equivalent according to the complete relation are also equivalent according to the Minimal relation.

7.3.4.8 Types with mutual dependencies on other types

The XTYPES type system includes types that have mutual dependencies on other types. These types are used to express “recursive” data structures such as trees. For example:

```
struct NodeData {
    long l_data;
};

struct TreeNode;

struct TreeNode {
    NodeData data;
    sequence<@external TreeNode> children;
};
```

More complex dependency cycles are possible where one type depends on another, which depends on another forming a dependency chain that eventually points back to the original type.

The “simple” algorithm to compute the `TypeIdentifier` of a type based on a hash of its `TypeObject` fails when types have mutual dependencies on each other because the construction of the `TypeObject` requires knowledge of the `TypeIdentifier` of all the dependent types, creating a circular dependency.

7.3.4.8.1 Background: Basic graph theory

The problem of types with mutual dependencies can be formulated in terms of directed graphs (digraphs). Given a set of types we define the “Type Dependency” digraph for those types as follows:

- The vertices in the graph are the types.
- The edges in the graph represent the direct dependencies between types, that is, if type T1 directly references type T2 (e.g. T1 is a structure and T2 is the type of a member, or T1 is a collection, and T2 is the type of the collection element).

A “directed path” in a digraph is a sequence of vertices where each vertex is connected to the next by a directed edge.

A “directed cycle” is a directed path that starts and ends on the same vertex.

Reachability relation: A vertex V1 is reachable from vertex V2 in the digraph if and only if there is a directed path from V2 to V1.

Strong connectivity relation: Two vertices V1 and V2 are **strongly connected** if and only if they are mutually reachable, that is, V1 is reachable from V2 and V2 is also reachable from V1.

Strong connectivity is an equivalence relation. The resulting partitions are called **Strongly Connected Components**.

The **kernel DAG** is defined as the digraph created by “combining” strongly connected components into a single vertex:

- Kernel DAG vertices: The strongly connected components
- Kernel DAG edges: There is an edge from a strongly connected component SCC1 to a strongly connected component SCC2 if and only if the original digraph contains some vertex belonging to SCC1 with an edge to a vertex belonging to SCC2.

A basic theorem in graph theory proves that Kernel DAG is acyclic, hence the name DAG which stands for Directed Acyclic Graph.

Figure 21 below shows an example digraph, its strongly connected components, and the corresponding Kernel DAG.




Figure 21 – Directed graph, Strongly Connected Components, and Kernel DAG

The strongly connectivity relation partitions the vertices in a digraph into subsets called **strongly connected components**. This is shown on the left part of the figure. The right side shows the Kernel DAG constructed using the strongly connected components as vertices. It is always a directed acyclic graph (DAG).

7.3.4.9 Computation of Type identifiers for types with mutual dependencies

7.3.4.9.1 Introduction

Mutual dependencies between types appear as directed cycles in the type dependency digraph. For example, the type dependency graph for the “tree” types declared above has a directed cycle involving the vertices “TreeNode” and “sequence<TreeNode>”. This is shown in Figure 22 below.




Figure 22 – Dependency graph derived from a set of type definitions

Type representation and type dependencies operate on the equivalence classes defined by the COMPLETE and MINIMAL type relations defined in Clause 7.3.4.7. Types belonging to the same equivalence class have the same TypeObject so they are treated as “the same type”. Depending on the relation (MINIMAL or COMPLETE) selected we will end up with a different set of types and type dependencies.

The algorithm to generate the TypeObjects and TypeIdentifiers is the same regardless of the equivalence relation chosen. To generate both the algorithm will be run two times, one for each equivalence relation.

The “basic” algorithm to compute `Hash TypeIdentifier` consists of hashing the serialized `TypeObject`. The construction of a `TypeObject` requires having the `TypeIdentifiers` of all the types the `TypeObject` depends on. Therefore this “basic” algorithm can handle only situations where the dependency graph does not have cycles, that is, it is a DAG.

The following clause defines a more general algorithm to construct `TypeIdentifiers` and `TypeObjects` that can also handle cycles in the dependency graph.

7.3.4.9.2 Algorithm

Let **EK** be the desired equivalence kind. Either `EK_COMPLETE` or `EK_MINIMAL`, which selects whether we are constructing the `TypeObjects` and `TypeIdentifiers` according to the `MINIMAL` or the `COMPLETE` equivalence relation.

Let **Types(EK)** a self-contained set of types (i.e. type equivalence classes) for the selected equivalence relation `EK`. By self-contained we mean a set of types that does not depend on any type outside the set.

Let `T` be an element of **Types(EK)** whose `TypeObject` and `TypeIdentifier` we wish to compute. The algorithm will construct the `TypeObject` and `TypeIdentifier` for all types in **Types(EK)** but it can be started with any type as an entry point.

1. Let **TypeDependencyDG(T)** be the dependency digraph that contains only the types that are reachable from `T`. If this graph has no cycles, then `T` is not affected by mutual dependencies and the `TypeIdentifier` can use the regular algorithm of hashing the serialized `TypeObjects`, which can be constructed recursively. Otherwise proceed to step 2.
2. Let **ReducedDependencyDG(T)** be the subdigraph of **TypeDependencyDG(T)** where all the vertices that have no outgoing edges are removed. These represent types that do not depend on any other types so their `TypeIdentifier` (and `TypeObject`) can be computed directly.
3. Identify the Strongly Connected Components of the **ReducedDependencyDG(T)**. Let **DependencyKernelDAG(T)** be the Kernel DAG of **ReducedDependencyDG(T)**.
4. Use a depth-first algorithm to compute the `TypeIdentifier` of the types on each Strongly Connected Component in **DependencyKernelDAG(T)**:
 - a. If the Strongly Connected Component (**SCC**) has a single type, then use the regular algorithm to compute its `TypeIdentifier` based on the type identifiers of all types it depends on. The depth first order ensures that those identifiers have already been computed.
 - b. If the Strongly Connected Component (**SCC**) has multiple types, then sort them using the lexicographic order of their fully qualified type name. Let **SCCIndex(U)** be the sort index of each type `U` belonging to the **SCC** starting with index 1 for the first type. For anonymous types concatenate the fully-qualified name of the containing type with the member name using “.” as the separator, for example “`MyModule::MyStruct.myMember`”.

- i. Temporarily set the `TypeIdentifier` of each `U` belonging to the SC to:
 - `discriminator = TI_STRONGLY_CONNECTED_COMPONENT`
 - `sc_component_id = {discriminator=EK, hash= 0}`
 - `scc_length = Number of types in SCC`
 - `scc_index = SCCIndex(U)` . Note that $1 \leq \text{scc_index} \leq \text{scc_length}$
- ii. Construct the `TypeObject` of all the types in the SC using the temporary `TypeIdentifier` for references to other types in the SCC. The depth first order ensures that `TypeIdentifier` for other types that the SCC depends on have already been computed.
- c. Place computed `TypeObjects` from step 4.b into a sequence `TypeObjectSeq` in the order of their `scc_index`.
- d. Serialize the `TypeObjectSeq` using the XCDR serialization for sequences with encoding version 2 and little endian.
- e. Compute the MD5 hash of the serialized buffer. Let **EquivalenceHash(SC)** be the first 14 bytes. Construct **StronglyConnectedComponentId(SC)** as:
 - i. `sc_component_id = { discriminator = EK, hash= EquivalenceHash(SC) }`
 - ii. `scc_length = Number of types in SCC`
- f. Set the `TypeIdentifier` of each of the types in SC to:
 - `discriminator = TI_STRONGLY_CONNECTED_COMPONENT`
 - `strong_component_id = StronglyConnectedComponentId(SC)`
 - `scc_index = SCCIndex(U)`

Implementation notes: (non-normative):

- The strongly connected component of a vertex `V` can be constructed as the set of vertices `W` reachable from `V` both by backwards and forwards traversal. If we define `Forward(V)` as the vertices reachable from `V` and `Backward(V)` as the set of vertices from which it is possible to reach `V`. Then:
 - `StronglyConnectedComponent(V) = Forward (V) ∩ Backward (V)`.
 - `Forward (V)` can be computed using depth first search (DFS) from `V`.
 - `Backward (V)` can be computed using DFS on the transpose graph obtained by inverting every edge.
- There are simple linear time algorithms (e.g. Kosaraju-Sharir) that compute the strongly connected components of a graph.

7.3.4.9.3 Strongly Connected Components Identifier (SCCIdentifier)

Each Strongly Connected Component (SCC) is uniquely identified by a **StronglyConnectedComponentId**. The **StronglyConnectedComponentId** is constructed using the algorithm specified in 7.3.4.9.2.

The **StronglyConnectedComponentId** contains the number of types in the strongly connected component (field *scc_length*) and a hash of all the corresponding `TypeObjects` (field *sc_component_id*).

From the **StronglyConnectedComponentId** it is possible to derive the `TypeIdentifiers` of all the types in the SCC. The `TypeIdentifiers` of all the types belonging to the same SCC only differ on the *scc_index* field, which always takes values from 1 to *scc_length*.

There are situations where an SCC needs to be identified without referencing a concrete type inside the SCC. In this situation a `TypeIdentifier` is constructed the same way as for any of the types in the SCC except the *scc_index* field is set to 0. We refer to this special `TypeIdentifier` recognizable by its discriminator being equal to `TI_STRONGLY_CONNECTED_COMPONENT` and *scc_index* = 0 AS the **SCCIdentifier**.

The `TypeIdentifier` of any type in the SCC contains the information needed to construct the **SCCIdentifier**.

7.4 Data Representation

The Data Representation module specifies the ways in which a data object of a given type can be externalized so that it can be stored in a file or communicated over the network. This is also commonly referred as “data serialization” or “data marshaling.”

Data Representations serve multiple purposes such as:

- Represent data in a “byte stream” so it can be sent over the network
- Represent data in a “byte stream” so it can be stored in a file
- Represent data in a human-readable form so it can be displayed to the user
- Provide a language for the user to enter data-values to a tool or specify them in a file

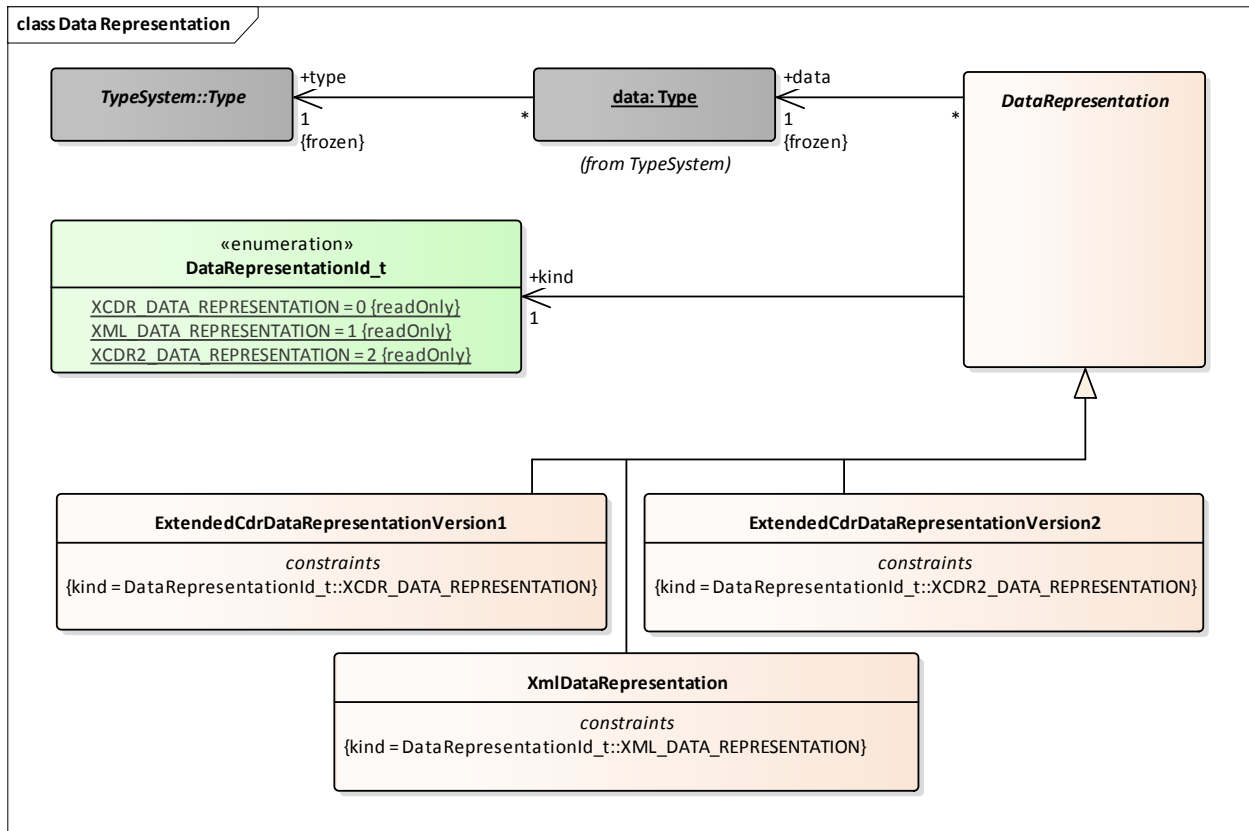


Figure 23 – Data Representation—conceptual model

This specification introduces multiple Data Representations. The reason for defining multiple type representations is that each of these is better suited or optimized for a particular purpose. These representations are all mostly equivalent. Consequently, other than convenience or performance, there is little reason to use one versus the other.

The alternative representations are summarized in Table 30.

Table 30 – Alternative Data Representations

<i>Data Representation</i>	<i>Reasons for using it</i>	<i>Disadvantages</i>
Extended CDR, encompassing both “traditional” CDR and parameterized CDR	<p>Compact and efficient binary representation. Minimizes CPU and Bandwidth used.</p> <p>Supports type evolution.</p> <p>Existing international OMG Standard. (Traditional CDR from CORBA [CDR]; parameterized CDR from RTPS [RTPS].)</p> <p>Already in used in the DDS Interoperability Protocol.</p>	Not human readable.
XML	<p>Human Readable</p> <p>Easily parsed and transformed with standard tools</p>	<p>CPU Intensive</p> <p>Uses 10 or 20 times more space than CDR</p>

7.4.1 Extended CDR Representation (encoding version 1)

This specification defines extensions of the OMG CDR representation [CDR] able to accommodate both optional members and appendable/mutable types. These extensions result in two encoding formats: PLAIN_CDR and PL_CDR.

Both encoding formats leverage the OMG CDR representation for all primitive types and non-mutable constructed types where the (traditional) CDR representation is well defined:

- PLAIN_CDR introduces extensions to CDR in order to handle optional members, bitmasks, and maps.
- PL_CDR leverages the RTPS Parameter List representation to handle mutable types.

7.4.1.1 PLAIN_CDR Representation

The PLAIN_CDR representation shall be used for final and appendable types, including (trivially) primitive types. It shall also be used for all string, sequence, and map types. Aggregated types declared as mutable shall use the PL_CDR representation described in Clause 7.4.1.2.

The PLAIN_CDR representation is based on the traditional CDR representation format [CDR] with the minimal extensions described below needed to handle the new types and concepts introduced by this specification.

The [RTPS] specification states that following the serialized data submessage element, padding bytes shall be added so that the next submessage starts at a 4-byte offset relative to the beginning of the RTPS message. This XTYPES specification further requires that any padding bytes added at the end of the serialized data shall be set to zero.

7.4.1.1.1 Primitive types

The PLAIN_CDR representation for primitive types shall be the same as in “traditional” CDR [CDR]. Specifically:

- The serialized data shall be encoded at an offset that aligned to the size of the primitive type.
- An endianness byte swap shall be performed in case the native system endianness is different from the one currently configured in the XCDR stream (XCDR.cendien).

Table 31 below summarizes the serialization of various primitive types.

Table 31 – Serialization of primitive types in version 1 encoding

<i>Primitive Type</i>	<i>Encoded Size</i>	<i>Alignment (version 1)</i>	<i>Byte representation</i>
Byte	1	1	The byte value
Boolean	1	1	0 for false, 1 for true
Char8	1	1	The character value encoded as described in 7.2.2.2.1.2
Char16	2	2	The character value encoded as described in 7.2.2.2.1.2
Int16 UInt16	2	2	The integer value using two’s complement notation
Int32 UInt32	4	4	The integer value using two’s complement notation
Int64 UInt64	8	8	The integer value using two’s complement notation
Float32	4	4	IEEE standard for normalized single-precision floating-point numbers [IEEE-748]
Float64	8	8	IEEE standard for normalized double-precision floating-point numbers [IEEE-748]
Float128	16	8	IEEE standard for normalized quadruple-precision floating-point numbers [IEEE-748]

7.4.1.1.2 Character Data

Objects of Char8 type shall not be interpreted to have a specific encoding and shall be serialized as-is in the same way as the Byte primitive type.

Objects of String<Char8> type shall be represented using the UTF-8 character encoding. The serialized length of an object of type String<Char8> shall be the number of bytes in the CDR buffer taken by the String<Char8> characters, including the terminating NUL character. The serialized length may not be the same as the number of Unicode characters because a single Unicode character encoded using the UTF-8 encoding may take one to four bytes.

Objects of `String<Char16>` type shall be represented using the UTF-16 character encoding. The serialized length of an object of type `String<Char16>` shall be the number of bytes in the CDR buffer taken by the `String<Char16>` characters. This is twice the number of characters in the string because a single character (in the Basic Multilingual Plane) encoded using UTF-16 takes 2 bytes to serialize.

The UTF-16 representation of object of type `String<Char16>` shall not include a Byte Order Mark (BOM). The representation shall also not include any terminating NUL character(s).

Rationale: By setting the serialized length equal to the number of bytes the representation could support sending UTF-16 encoded Unicode characters outside the BMP (which map to two UTF-16 units). In this case, the serialized length would still indicate the number of bytes until the end of the string. The byte order used by the UTF-16 representation can be inferred from the one already available in the RTPS Encapsulation Identifier (see Clause 7.6.2.1.2), therefore the BOM is not needed. Finally terminating UTF-16 encoded strings with NUL characters is not considered best practice and the latest versions of OMG CDR do not do it.

7.4.1.1.3 Enumerated Types

7.4.1.1.3.1 Enumeration Types

Objects of enumerated types shall be serialized as integers, the sizes of which shall depend on the “bit bound” of their associated type.

Table 32 – Serialization of enumeration types

<i>Corresponding Primitive Type</i>	<i>Bit Bound</i>
Byte	1-8
Int16	9-16
Int32	17-32 (32 bits is the default size, and corresponds to all enumerated types prior to this specification)

7.4.1.1.3.2 Bitmask Types

Objects of bitmask types shall be serialized in the same way as the following primitive types, depending on the bitmask’s bound:

Table 33 – Serialization of bitmask types

<i>Bound</i>	<i>Corresponding Primitive Type</i>
[1..8]	Byte
[9..16]	UInt16
[17..32]	UInt32
[33..64]	UInt64

Bit indexes are counted from zero starting at the least-significant bit of the full byte size of the bitmask. In the case where the bound of the bitmask is less than the number of bits in the corresponding primitive type, the states of the remaining serialized bits are not specified, and those bits are not considered to be part of the bitmask.

7.4.1.1.4 Map Types

Objects of map types shall be represented according to the following equivalent IDL:

```
struct MapEntry_<key_type>_<value_type>[_<bound>] {
    <key_type> key;
    <value_type> value;
};
```

```
typedef sequence<MapEntry_<key_type>_<value_type>[_<bound>][, <bound>]>
Map_<key_type>_<value_type>[_<bound>];
```

The *<key_type>* and *<value_type>* names are as defined the Type System. See also Clause 7.2.2.4.3, which defines the implicit names of collection types.

For example, objects of the following IDL map type:

```
map<long, float>
```

...shall be serialized as if they were of the following IDL sequence type:

```
struct MapEntry_Int32_Float32 {
    long key;
    float value;
};

typedef sequence<MapEntry_Int32_Float32> Map_Int32_Float32;
```

7.4.1.1.5 Structures

Objects of structure type shall be represented as defined by the CDR specification [CDR], augmented as described below.

7.4.1.1.5.1 Inheritance

The members defined by the base type, if any, shall be serialized before the members of their derived types. The representation shall be exactly as if all of the members had been defined, in the same order, in the most-derived type.

7.4.1.1.5.2 Optional Members

Structure members marked as optional shall be preceded by a parameter header as described in Clause 7.4.1.2, “Parameterized CDR Representation”, below.

7.4.1.2 Parameterized CDR Representation

The parameterized CDR representation is based on the RTPS Parameter List CDR representation defined in [RTPS].

Each element, or parameter, within a parameter list data structure is simply a CDR-encapsulated block of data. Preceding each one is a parameter header consisting of a two-byte parameter ID followed by a two-byte parameter length. One parameter follows another until a list-terminating sentinel is reached.

Unlike it is stated in [RTPS] Sub Clause 9.4.2.11 “ParameterList”, the value of the parameter length is the exact length of the serialized member. It does not account for any padding bytes that may follow the serialized member. Padding bytes may be added in order to start the next parameterID at a 4 byte offset relative to the previous parameterID.

This data representation uses elements of the parameter list data structure for two purposes:

- Any object of a mutable aggregated type shall be serialized as a parameter list. Each of its members shall correspond to a single parameter within that list.
- Any optional member of a final or appendable structure shall be preceded by a parameter header describing that member. If the member takes no value within that particular object, the data length indicated by the header shall be zero. This reuse of the parameter header data structure does not constitute a complete parameter list: the optional member shall not be followed by list-terminating sentinel.

7.4.1.2.1 Interpretation of Parameter ID Values

As described in Clause 9.6.2.2.1, *ParameterId space*, of the RTPS Specification, the 16-bit-wide parameter ID range may be interpreted as a two-bit-wide bitmask followed by a 14-bit wide unsigned integer.

- The first bit of the bitmask—the most-significant bit of 16-bit-wide the parameter ID as a whole—indicates whether the parameter has an implementation-specific interpretation. This specification refers to this bit as `FLAG_IMPL_EXTENSION`.
- The second bit of the bitmask indicates whether the parameter, if its ID is not recognized by the consuming implementation, may be simply ignored or whether it causes the entire data sample to be discarded. This specification refers to this bit as `FLAG_MUST_UNDERSTAND`. This bit shall be set if and only if the `must_understand` property of the member being encapsulated is set to true.

Within the 14-bit-wide integer region of the parameter ID, this specification further reserves the largest 255 values—from 16,129 (0x3F01) to 16,383 (0x3FFF)—for use by the OMG in this specification and future specifications. **Table 34** below identifies the reserved parameter ID values. For a parameter to be recognized as one of the well-known values in **Table 34**, the `FLAG_IMPL_EXTENSION` bit must be set to zero. Refer to **Table 34** for the value of the `FLAG_MUST_UNDERSTAND` bit.

Table 34 – Reserved parameter ID values

<i>Name</i>	<i>14-Bit Hex Value(s)</i>	<i>FLAG_MUST_UNDERSTAND set?</i>	<i>Description</i>
PID_EXTENDED	0x3F01	Yes	Allows the specification of large member ID and/or data length values; see below
PID_LIST_END	0x3F02	Yes	Indicates the end of the parameter list data structure. RTPS specifies that the PID value 1 shall be used to terminate parameter lists within the DDS built-in topic data types. Rather than reserving this parameter ID for all types, thereby complicating the member ID-to-parameter ID mapping rules for all producers and consumers of this Data Representation, <i>Simple Discovery types</i> shall be subject to a special limitation: member ID 1 shall not be used and parameter ID 1 shall terminate the parameter list to provide backwards compatibility. Implementations shall be robust to receiving parameter ID 0x3F02 to indicate the end of a list as well. These types consist of the built-in topic data types, and those other types that contain them as members, as defined by [RTPS].
PID_IGNORE ⁵	0x3F03	No	All consumers of this Data Representation shall ignore parameters with this ID.
<i>Reserved for OMG</i>	0x3F04–0x3FFF	<i>N/A</i>	<i>Reserved for OMG</i>

When writing data, implementations of this specification shall set the `FLAG_MUST_UNDERSTAND` bit as described in Table 34. When reading data, implementations of this specification shall be robust to any setting of the `FLAG_MUST_UNDERSTAND` bit and accept the parameter nevertheless.

⁵ **Design rationale (non-normative):** RTPS uses PID 0 (“PID_PAD”), corresponding to member ID 0, as a padding field. PID_IGNORE applies this concept to all data types using this Data Representation. The additional reservation of PID 0 is not necessary: because the types defined by RTPS do not use member ID 0, consumers of those types will naturally ignore any incidence of its corresponding PID that they encounter.

This specification extends the parameter list data structure to permit 32-bit parameter IDs and data lengths up to 4 Giga-Bytes. This extension uses the reserved must-understand 16-bit parameter ID `PID_EXTENDED` to indicate that a member's parameter ID and/or length require 32-bits. The member ID (long member ID) and member length (long member length) follow in the 8 bytes directly after the `PID_EXTENDED` parameter ID and accompanying 16-bit length.

The value of the `PID_EXTENDED` with the must understand flag set is `0x7F01` (that is `0x4000 + 0x3F01`).

The four bytes following the `PID_EXTENDED` and length shall be a serialized `UINT32` value "eMemberHeader" that is constructed by combining four 1-bit flags with by the 28-bit member ID. The flags occupy the 4 most significant bits of the `UINT32` value. The flags are combined with the `memberId` as shown below:

```

FLAG_1 = 0x80000000
FLAG_2 = 0x40000000
FLAG_3 = 0x20000000
FLAG_4 = 0x10000000
eMemberHeader = FLAG_1 + FLAG_2 + FLAG_3 + FLAG_4 + memberId

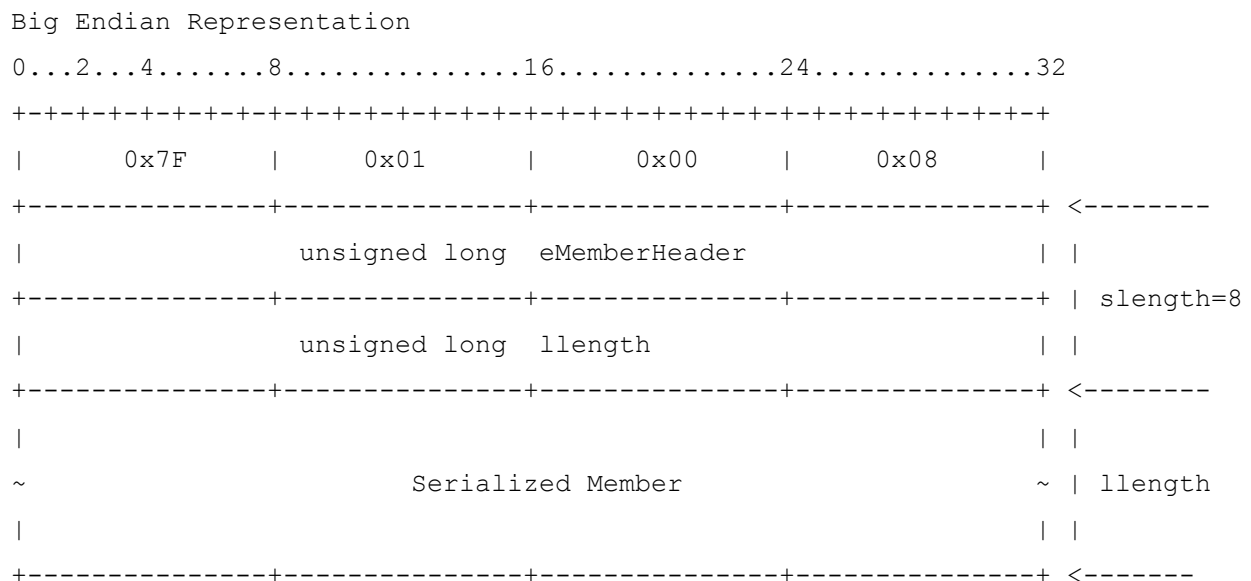
```

The second four bytes following the `PID_EXTENDED` and length shall be interpreted as a 32-bit unsigned integer (`llength`) that contains the length of the serialized member. Note that `llength` is the exact length of the serialized member and does not account for any padding that may follow the member.

The value of the 16-bit length associated with the `PID_EXTENDED` (`slength`) shall be equal to eight.

The serialized member shall start immediately after the long member length (`llength`). That is exactly 12 bytes from the beginning of the `PID_EXTENDED` parameter.

See Figure 24 for an example of the layout of the CDR buffer where `PID_EXTENDED` is used.






Figure 24 – Usage of PID_EXTENDED within the CDR Buffer

The setting of the `FLAG_IMPL_EXTENSION` and `FLAG_MUST_UNDERSTAND` bits in the 16-bit parameter ID shall not be interpreted to apply to the extended parameter as well. Instead, the first most-significant bit of the four-bitmask of flags within the extended parameter header shall represent the value of `FLAG_IMPL_EXTENSION` for the data member. The second most-significant bit shall represent the `FLAG_MUST_UNDERSTAND` value of the data member. The remaining two bits, unless specified by some other OMG specification, should be set to zero.

These extended parameter headers, based on `PID_EXTENDED`, shall be legal within the parameter list data structures used to serialize objects of mutable aggregated types. They shall also be legal when preceding optional members of final or appendable structures, as described above.

The alignment rules for extended parameters shall be the same as those for non-extended parameters, which are defined in [RTPS] Clause 9.4.2.11.

7.4.1.2.2 Member ID-to-Parameter ID Mapping

The mapping from member IDs to parameters shall be as follows:

- Member IDs from 0 to 16,128 (0x3F00) inclusive shall be represented exactly in the lower 14 bits of the parameter ID.
- All other member IDs must be expressed using the extended parameter header format.
- Almost any parameter can legally be expressed using extended parameter headers. There is no requirement that parameters that *could* be described with the shorter header defined by the RTPS Specification *must* be described that way; if a parameter could be described using a short parameter header or an extended header, the short and extended expressions of that header shall be considered totally equivalent. This mapping ensures that members of user-defined data types will never set the `FLAG_IMPL_EXTENSION` bit. Currently, the

`FLAG_IMPL_EXTENSION` bit is used only for RTPS discovery-defined data types, which may or may not have the bitmask as defined by the RTPS Specification itself.

7.4.1.2.3 Omission and Reordering of Members of Aggregated Types

Because each parameter (type member, in this case) is explicitly identified, and identification of mutable structure members occurs based on the IDs of those parameters, members of mutable structures may appear in any order. Furthermore, any mutable structure member's value may be omitted. In such a case, if the member is not optional, it logically takes its default value. If the member is optional, it takes no value at all.

Objects of final or appendable structures are not serialized as full parameter lists, even if some members are optional. Therefore, the members of these types may not be omitted or reordered.

Because union members are identified based on a discriminator value, the value of the discriminator member must be serialized before the value of the current non-discriminator member. Neither member value may be omitted.

7.4.1.2.4 Nested Objects

In the case where an object of an aggregated mutable type contains another object of an aggregated mutable type, one parameter list will contain another. In that case, parameter IDs are interpreted relative to the innermost type definition. (For instance, a type `Foo` may contain an instance of type `Bar`. Both `Foo` and `Bar` may define a member with ID 5. Inside the parameter list corresponding to the `Bar` object, an occurrence of parameter ID 5 shall be considered to refer to `Bar`'s member 5, not to `Foo`'s member 5.)

Likewise, an occurrence of `PID_LIST_END` indicates the conclusion of the innermost parameter list.

7.4.2 Extended CDR Representation (encoding version 2)

This specification defines three encoding formats used with encoding version 2: `PLAIN_CDR2`, `DELIMITED_CDR`, and `PL_CDR2`.

The three encoding formats leverage the `PLAIN_CDR` representation. They enhance the encodings used in version 1 to improve type assignability and reduce the size of serialized data.

- `PLAIN_CDR2` shall be used for all primitive, strings, and enumerated types. It is also used for any type with extensibility kind `FINAL`. The encoding is similar to `PLAIN_CDR` except that `INT64`, `UINT64`, `FLOAT64`, and `FLOAT128` are serialized into the CDR buffer at offsets that are aligned to 4 rather than 8 as was the case in `PLAIN_CDR`.
- `DELIMITED_CDR` shall be used for types with extensibility kind `APPENDABLE`. It serializes a `UINT32` delimiter header (`DHEADER`) before serializing the object using `PLAIN_CDR2`. The delimiter encodes the endianness and the length of the serialized object that follows.
- `PL_CDR2` shall be used for **aggregated types** with extensibility kind `MUTABLE`. Similar to `DELIMITED_CDR` it also serializes a `DHEADER` before serializing the

object. In addition it serializes a member header (EMHEADER) ahead each serialized member. The member header encodes the member ID, the must-understand flag, and length of the serialized member that follows.

7.4.3 Extended CDR encoding virtual machine

The encoding formats are specified using a virtual machine that acts on an XCDR stream object. The XCDR stream holds the bytes resulting from the incremental serialization of data objects into the stream.

The XCDR stream model consists of:

- A linear byte buffer where the serialized objects are placed.
- A set of internal state variables that may affect the serialization of future objects serialized into the stream. See Table 36.
- A set of operations on the stream that modify the state variables. See Table 37.
- A “stream insertion” operation that serializes objects onto the stream with a format that depends on the object type, its composition, and the value of the state variables. The append operation is represented using the operator symbol “<<”. See Table 37.

7.4.3.1 Encoding version and format

The encoding format is determined by the encoding version and the extensibility kind of the object being serialized. Table 35 specifies the format that shall be used in each case.

Table 35 – Serialization format to use.

<i>Extensibility Kind</i>	<i>Encoding Version</i>	<i>Encoding format on the wire</i>
FINAL	1	PLAIN_CDR
FINAL	2	PLAIN_CDR2
APPENDABLE	1	PLAIN_CDR
APPENDABLE	2	DELIMITED_CDR
MUTABLE	1	PL_CDR
MUTABLE	2	PL_CDR2

7.4.3.2 XCDR Stream State

7.4.3.2.1 XCDR stream state variables

The state of the XCDR stream is described by the value of the variables (the XCDR state variables) defined in Table 36.

Table 36 – State variables and constants in the XCDR stream model

<i>XCDR state variable</i>	<i>meaning</i>
NENDIAN	<p>Constant that represents the native endianness used by the system. It is dependent on the processor architecture, compiler, and operating system.</p> <p>There are two possible values: LITTLE_ENDIAN and BIG_ENDIAN</p>
cendian	<p>Choice variable representing the current endianness. This is the endianness that will be used to serialize subsequent objects into the stream. It affects integer types, floating-point types, enumerated types, and the Char16 type.</p>
offset	<p>Integer variable representing the offset into the byte stream where the next serialized byte will be placed.</p> <p>XCDR.offset is computed relative to the beginning of the stream so that XCDR.offset counts the number of bytes currently serialized into the stream.</p> <p>Each byte serialized into the stream causes XCDR.offset to be incremented.</p>
origin	<p>Integer state variable representing the offset into the stream used as the “logical beginning of the stream” for alignment operations.</p> <p>Each Type “T” has a default alignment (T.dalignment). This is the alignment used by default when an object of that type is serialized into a stream.</p> <p>An object O of type T shall be serialized at an offset that verifies:</p> $((\text{XCDR.offset} - \text{XCDR.origin}) \% \text{T.dalignment}) == 0$ <p>If the current XCDR.offset does not satisfy the above condition, the serialization shall insert the minimum “padding bytes” needed to advance XCDR.offset so that the condition is met.</p>
eversion	<p>Octet state variable used to identify the version of the encoding rules used to serialize the stream.</p> <p>The pre-defined values are:</p> <p>{0x00} -- VERSION_NONE</p> <p>{0x01} -- VERSION1</p> <p>{0x02} -- VERSION2</p>

maxalign	<p>Integer state variable representing the maximum value for the alignment that will be used for future objects serialized into the stream. This value overrides the required alignment for the object being serialized, so the alignment condition for any object O of type O.type becomes:</p> $((XCDR.offset - XCDR.origin) \% MALIGN(O)) == 0$ <p>Where</p> $MALIGN(O) = \text{MIN}(O.type.alignment, XCDR.maxalign)$ <p>This value is automatically set from the XCDR.eversion.</p> $XCDR.maxalign == \text{MAXALIGN}(XCDR.eversion)$
----------	--

7.4.3.2.2 Operations that change the XCDR stream state

The XCDR stream state is modified as a result of the serialization of data objects into the stream. It can also be modified as a result of performing the operations shown in Table 37.

Table 37 – Stream operations in the XCDR stream model

<i>XCDR stream operation</i>	<i>meaning</i>
INIT(V1=<nv1>, V2=<nv2>,...)	<p>Initializes (constructs) the XCDR stream and sets the state variables V1, V2, ... as specified.</p> <p>The notation <?> indicates that the value can be chosen by the implementation.</p>
PUSH(VARIABLE=<newvalue>)	<p>Pushes the specified XCDR stream variable VARIABLE into the stack and sets the current value to <newvalue>.</p> <p>The notation <?> indicates that the new value can be chosen by the implementation.</p> <p>This action is reverted by the POP() operation.</p>
PUSH(V1=<nv1>,V2=<nv2>,...)	<p>A shortcut for calling PUSH() multiple times with the listed variables and new values.</p>
POP(VARIABLE)	<p>Replaces the XCDR stream variable VARIABLE with the value for that variable that was pushed on the last PUSH() operation, removing it from the stack.</p>
POP(V1, V2,...)	<p>A shortcut for calling POP() multiple times with the listed variables.</p>

MAXALIGN(<eversion>)	<p>This operation returns the maximum alignment used for a given version of the encoding:</p> <p>MAXALIGN(VERSION2) = 4</p> <p>MAXALIGN(VERSION1) = 8</p> <p>MAXALIGN(VERSION_NONE) = 8</p>
ALIGN(N)	<p>This operation is used to advance the XCDR stream to achieve a desired alignment of the XCDR.offset.</p> <p>Advancing the XCDR.offset is done by inserting “padding bytes” into the stream. The value of the padded bytes is left unspecified.</p> <p>The actual number of bytes advanced depends not only on “N” but also on the value of the XCDR.maxalign. Specifically the stream is aligned to neededalign:</p> <p>neededalign = MIN(N, XCDR.maxalign)</p> <p>After the operation is performed the following condition shall be true:</p> <p>$(\text{XCDR.offset} - \text{XCDR.origin}) \% \text{neededalign} == 0$</p>
XCDR << { O : T }	<p>The “append” stream operation.</p> <p>Serializes (using the Extended CDR representation) an object “O” of type “T” onto the XCDR stream starting at offset XCDR.offset.</p>

7.4.3.2.3 XCDR Stream Initialization

The XCDR stream shall be initialized with an empty buffer.

The endianness shall be set as desired by the implementation, although a common setting for best performance is the native system endianness (NENDIAN).

The encoding version (eversion) shall be set as configured on the DataWriter. In this version of the DDS-XTypes specification it may be set to 1 or 2.

The first 2 octets in the XCDR stream shall be the Encapsulation Header (ENC_HEADER) indicating the endianness, encoding version, and encoding algorithm of the top-level type. See Table 39. This is the type associated with the DataWriter.

7.4.3.3 Type and Byte transformations

The operation of the serialization virtual machine uses a set of helper type and byte-buffer transformations.

The type transformations transform a type into another type, typically modifying its extensibility kind.

The byte-buffer transformations perform byte swaps in arrays of bytes or allow reinterpreting an object of a primitive type as an array of bytes.

These transformations are used to decompose the serialization of one type as a set of serializations of other types which have already been described.

Table 38 defines the type and byte transformations.

Table 38 – Type and Byte transformations used in the serialization virtual machine

<i>Type or Object transformation</i>	<i>meaning</i>
AsFinal(T) for any type T	<p>This transformation only affects Aggregated types. For other types AsFinal(T) returns T.</p> <p>For the affected types AsFinal(T) is a new type which is declared the same as T except that its extensibility kind is FINAL.</p>
AsNested(T) for any type T	<p>This transformation treats the type as a Nested type for serialization purposes.</p>
AsBytes(O) for any object O of a PRIMITIVE_TYPE	<p>This transformation reinterprets the primitive object as an array of bytes.</p> <p>The resulting bytes are ordered as they appear in the processor memory according to the native Endianness (NENDIAN) used by the system.</p>
ESWAP(B, <doit>) where B is a stream of 1, 2, 4, or 8 bytes	<p>Conditionally swaps the bytes on the input stream B based on whether the current XCDR endianness (XCDR.cendian) matches the native Endianness (NENDIAN).</p> <p>This operation returns the same input stream if the input is a single byte or if XCDR.cendian == NENDIAN.</p> <p>Otherwise the operation produces a new stream of bytes with the same length as the input performing an (endianness) byte swapping according to the length of the input stream:</p> <p>For length 2: { B[1], B[0] }</p> <p>For length 4: { B[3], B[2], B[1], B[0] }</p> <p>For length 8: { B[7], B[6], B[5], B[4], B[3], B[2], B[1], B[0] }</p>

7.4.3.4 Functions related to data types and objects

The operation of the serialization virtual machine uses a set of helper functions that return bytes or data to append to the XCDR stream. The notation and meaning is defined in Table 39.

Table 39 – Functions operating on objects and types

<i>function</i>	<i>meaning</i>
ENC_HEADER(<E>, <eversion>, T) for any type “T”	ENC_HEADER is an array of 2 octets used to identify the type of encoding (serialization), version of the encoding (<eversion>) and the endianness used by the stream (<E>): {0x00, 0x00} -- PLAIN_CDR, BIG_ENDIAN, {0x00, 0x01} -- PLAIN_CDR, LITTLE_ENDIAN {0x00, 0x02} -- PL_CDR, BIG_ENDIAN, {0x00, 0x03} -- PL_CDR, LITTLE_ENDIAN, {0x00, 0x10} -- PLAIN_CDR2, BIG_ENDIAN, {0x00, 0x11} -- PLAIN_CDR2, LITTLE_ENDIAN {0x00, 0x12} -- PL_CDR2, BIG_ENDIAN {0x00, 0x13} -- PL_CDR2, LITTLE_ENDIAN {0x00, 0x14} -- DELIMIT_CDR, BIG_ENDIAN {0x00, 0x15} -- DELIMIT_CDR, LITTLE_ENDIAN {0x01, 0x00} -- XML
EVERSION(T) for any type “T”	EVERSION is an octet used to identify the version of the encoding rules used to serialize the stream. The values are: 0x00 -- unspecified version (understood as version 1) 0x01 -- version 1 0x02 -- version 2

<p>DHEADER(O, E) for any object O of type T</p>	<p>A UInt32 header value computed as the sum: $DHEADER(O) = (E_FLAG \ll 31) + O.size$ Where E is set as desired by the implementation: E = 1 indicates that following the header XCDR stream endianness shall be changed to LITTLE_ENDIAN. E = 0 indicates that following the header XCDR stream endianness shall be changed to BIG_ENDIAN. O.size is the number of bytes following the header that are required to hold the serialized representation of O.</p>
<p>EMHEADER1(M) Where M is a member of a structure</p>	<p>EMHEADER1 is the first 4 bytes of the Enhanced Mutable Header (EMHEADER) is used by the PL_CDR2 encoding format. It is a UInt32 value computed as: $EMHEADER1 = (M_FLAG \ll 31) + (LC \ll 28) + M.id$ Where: M_FLAG is the value of the Must Understand option for the member LC is the value of the Length Code for the member.</p>
<p>LC(M) Where M is a member of a structure</p>	<p>LC is a 3-bit length code used to construct the EMHEADER1. It determines whether EMHEADER header has an additional 4 bytes (the NEXTINT) and is also used to encode the serialized size of the member that follows.</p>
<p>NEXTINT(M) Where M is a member of a structure</p>	<p>NEXTINT is the second 4 bytes of the Enhanced Mutable Header (EMHEADER). It is a UInt32 value. NEXTINT is only present if $LC(M) \geq 4$. NEXTINT is used in combination with LC to encode the serialized size of the member that follows.</p>

7.4.3.4.1 Delimiter Header (DHEADER)

The DELIMITED_CDR and PL_CDR encoding formats prepend a UInt32 delimiter header (DHEADER) ahead of the serialization of the object content.

The DHEADER encodes the endianness used to serialize the object as well as the serialized size of the object that follows (not including the DHEADER itself). It is computed with the formula:

$$DHEADER(O) = (E_FLAG \ll 31) + (O.size \& 0x8fffffff)$$

In this expression, O.size is constrained to being smaller than 2 Giga Bytes (2³¹ Bytes) and E_FLAG is set to 0 if the object will be serialized using big endian serialization and 1 if it will use little endian.

The serialization of the DHEADER being a Uint32 type forces a 4-byte alignment relative to XCDR.origin, this may insert into the stream up to 3 padding bytes prior to the DHEADER.

The serialization of the DHEADER uses the endianness active in the XCDR stream at the time it is serialized (XCDR.cendian). Following the serialization of DHEADER the value of the endianness encoded into the header (E_FLAG) shall be pushed into the XCDR stream.

7.4.3.4.2 Member Header (EMHEADER), Length Code (LC) and NEXTINT

The PL_CDR2 encoding format serializes aggregated types using a member-by-member Type-Length encoding.

A member header precedes the serialization of each member. The member header can be either 4 or 8 bytes.

The first four bytes are the serialized representation of a Uint32 integer called EMHEADER1. EMHEADER1 shall be serialized using the XCDR stream endianness current at the place the serialization occurs (XCDR.cendian).

The second 4 bytes, if present, are the serialized representation of a Uint32 integer called NEXTINT. It shall be serialized with the same endianness as EMHEADER1.

EMHEADER1 is constructed from three parts: The must understand flag (M_FLAG), the length code (LC) and the member ID.

```
EMHEADER1 = (M_FLAG << 31) + (LC << 28) + (MemberId & 0x0fffffff)
```

The must understand flag (M_FLAG) shall be set to 1 if the corresponding member must be understood by the receiver, see Clause 7.2.2.4.4.4. Otherwise it shall be set to zero.

The length code provides the means to determine the serialized size of the member. There are eight possible values from 0 to 7 both included (0b000 to 0b111 in binary). These are interpreted as follows:

- LC values between 0 and 3 indicate that the member header is 4 bytes. That is, there is no NEXTINT. The value of LC encodes the length of the serialized member directly:
 - LC = 0 = 0b000 indicates serialized member length is 1 Byte
 - LC = 1 = 0b001 indicates serialized member length is 2 Bytes
 - LC = 2 = 0b010 indicates serialized member length is 4 Bytes
 - LC = 3 = 0b011 indicates serialized member length is 8 Bytes

- LC values between 4 and 7 indicate that the member header is 8 bytes. That is, a second integer (NEXTINT) immediately follows EMHEADER1. The value of LC combined with the value NEXTINT encode the length of the serialized member:
 - LC = 4 = 0b100 indicates serialized member length is NEXTINT
 - LC = 5 = 0b101 indicates serialized member length is also NEXTINT
 - LC = 6 = 0b110 indicates serialized member length is 2*NEXTINT
 - LC = 7 = 0b111 indicates serialized member length is 4*NEXTINT

EMHEADER1 with LC values 5 to 7 also affect the serialization/deserialization virtual machine in that they cause NEXTINT to be reused also as part of the serialized member. This is useful because the serialization of certain members also starts with an integer length, which would take exactly the same value as NEXTINT. Therefore the use of length codes 5 to 7 saves 4 bytes in the serialization.

7.4.3.5 Encoding (serialization) rules

The logic of the virtual machine is expressed as a collection of rules. Each rule has the form:

```
XCDR[vv] “<<” <match criteria> “=” XCDR “<<” <serialization action1>
“<<” <serialization action2>
“<<” ...
```

XCDR represents the stream containing the serialization of an object. It has a state represented by its state variables (see Clause 7.4.3.1) and it also holds the bytes from previously serialized objects. The [*vv*] indicates the encoding version that the DataWriter uses. This is configured on each DataWriter. A stream has its encoding version set when it is initialized and it cannot be modified.

A rule with left hand side XCDR[*vv*] only applies if the XCDR stream is using encoding version *vv*. A rule with left hand side XCDR applies for all xtypes encoding versions.

The **<match criteria>** represents the object that is being serialized into the XCDR stream.

When serializing an object each rule is evaluated in sequence and the first one that has a matching version and criteria is applied.

The application of a rule consists of executing each one of the serialization actions. Each action may change state variables of the stream or indicate that new objects (or modifications to existing objects) shall be serialized. This may recursively trigger the application of new rules.

The rules shall be applied until completion. Once completed, the XCDR stream contains the serialized representation of the object that initiated the serialization.

The rules are written from the point of view of a writer that is constructing the RTPS SerializedData buffer to send. Therefore the entrypoint is a so-called “Top Level” type which indicates a non-nested type that can be published by a DDS DataWriter. This entry point ensures the XCDR stream includes the SerializedData encapsulation header required by the DDS-RTPS

protocol. Other entry points are possible if the intent is to simply serialize an object and not embed it within an RTPS SerializedData.

7.4.3.5.1 Notation used for the match criteria

Table 40 shows the symbols and notation used by the serialization virtual machine.

Table 40 – Symbols and notation used in the serialization virtual machine

<i>notation</i>	<i>meaning</i>
O : T	An object “O” of type “T” <ul style="list-style-type: none"> • O.type is another way to refer to the object type “T” • O.ssize is the size in bytes required to hold the serialized representation of O in an XCDR stream that has XCDR.offset aligned to the T.dalignment.
O : TOP_LEVEL_TYPE	An object O being serialized as the top-level Topic-Type. That is as the object written directly by a data-writer and not a nested object.
O : PRIMITIVE_TYPE	An object O of a primitive type as defined in 7.2.2.2.
O : STRING_TYPE	An object O of a string type which Char8 elements as defined in 7.2.2.4.3
O : WSTRING_TYPE	An object O of a string type with Char16 elements as defined in 7.2.2.4.3
O : ENUM_TYPE	An object “O” of an Enumerated type as defined in 7.2.2.4.1 <ul style="list-style-type: none"> • O.holder_type is either Byte, Int16 or Int32 depending on the value of the @bit_bound annotation. • O.value is the (integer) value of the enumeration.
O : BITMASK_TYPE	An object O of a BitMask type as defined in 7.2.2.4.1.2 <ul style="list-style-type: none"> • O.holder_type is UInt16, UInt32, or UInt64 depending on the value of the @bit_bound annotation. • O.value is the (integer) value of the bitmask.
O : ALIAS_TYPE	An object O of an Alias type as defined in 7.2.2.4.2 <ul style="list-style-type: none"> • O.base_type is the equivalent (aliased) type.

O : ARRAY_TYPE	<p>An object “O” of an Array type as defined in 7.2.2.4.3</p> <ul style="list-style-type: none"> • O.element_type is the element type • O.length is the total number of elements in the array (accounting for all the dimensions) <p>For single- dimensional arrays O[i] is the “ith” element in the array.</p> <p>Multi-dimensional arrays are treated for serialization purposes as a single dimensional array containing all the elements ordered such that the index of the first dimension varies most slowly, and the index of the last dimension varies most quickly.</p>
O: FARRAY_TYPE	Same as ARRAY_TYPE except that its extensibility kind is FINAL.
O: PARRAY_TYPE	An ARRAY_TYPE whose element type is primitive.
O : SEQUENCE_TYPE	<p>An object “O” of a Sequence type as defined in 7.2.2.4.3</p> <ul style="list-style-type: none"> • O.element_type is the element type • O.length is the number of elements in the sequence. <p>Empty sequences have O.length==0</p> <p>For non empty sequences O[i] is the “ith” element in the sequence.</p> <p>Sequence indices are zero-based so O[0] is the first element in the sequence and O[O.length-1] is the last element in the sequence.</p>
O : PSEQUENCE_TYPE	<p>Same as SEQUENCE_TYPE except that O.element_type is a primitive type.</p> <p>These sequences are intrinsically delimited in the sense that the CDR representation allows determining the serialized size of the entire sequence without iterating over each element.</p>
O: FSEQUENCE_TYPE	Same as SEQUENCE_TYPE except that its extensibility kind is FINAL.

O : MAP_TYPE	<p>An object “O” of a Map type as defined in 7.2.2.4.3</p> <ul style="list-style-type: none"> • O.key_type is the key type • O.element_type is the element type • O.length is the number of keys in the map, which is also the number of elements in the map. <p>For non empty maps O[i].key is the “ith” key in the map, O[i].element is the (value) element that corresponds to that key.</p> <p>Map indices are zero-based so O[0].key is the first key in the map and O.key[O.length-1] is the last key in the map.</p>
O : FMAP_TYPE	A MAP_TYPE whose extensibility kind is FINAL.
O : PMAP_TYPE	A MAP_TYPE whose element and key are primitive types.
O : UNION_TYPE	<p>An object “O” of a Union type as defined in 7.2.2.4.4.2</p> <ul style="list-style-type: none"> • O.disc is the discriminator member. • O.disc.value is the value of the discriminator member. • O.disc.type is the type of the discriminator member. • O.selected_member is the member of the union selected based on the value of the discriminator. Note that certain discriminator values may select no member. • O.selected_member.value is the value of the selected member, if any. • O.selected_member.type is the type of the selected member.
O: FUNION_TYPE	Same as UNION_TYPE except that its extensibility kind is FINAL.

O : STRUCT_TYPE	<p>An object “O” of a Struct type as defined in 7.2.2.4.4.1</p> <ul style="list-style-type: none"> • O.base_type is the type of the base Structure in case O.type inherits from another structure. • O.member_count is the number of members. <p>For non empty structures:</p> <ul style="list-style-type: none"> • O.member[i] is the “ith” member in the structure. It is a holder for the object that contains the value of the member and contains additional information. • Member indices are zero-based so O[0] is the first member. <p>See definition of MEMBER.</p>
O : FSTRUCT_TYPE	Same as STRUCT_TYPE except that its extensibility kind is FINAL.
O : MSTRUCT_TYPE	<p>Same as STRUCT_TYPE except that its extensibility kind is MUTABLE.</p> <p>Unlike FSTRUCT_TYPE, O.member[i].id is the MemberId of O.member[i] as defined in 7.2.2.4.4.3 which may be different from “i”.</p>
M : MEMBER	<p>A member of an Aggregated type, 7.2.2.4.4.</p> <ul style="list-style-type: none"> • M.id is the member ID. • M.value is the object holding the value of the member. • M.value.type is the type of the object. • M.value.ssize is the serialized size of the object holding the value of the member.
M : FMEMBER	A member (see MEMBER) of an Aggregated type that has extensibility kind FINAL.
M : OPT_FMEMBER	A optional member (see Clause 7.2.2.4.4.5) of an Aggregated type with extensibility kind final (FMEMBER).
M : NOPT_FMEMBER	A non-optional member (see Clause 7.2.2.4.4.5) of an Aggregated type with extensibility kind final (FMEMBER).
M : MMEMBER	A member (see MEMBER) of an Aggregated type that has extensibility kind MUTABLE.
O : FINAL_TYPE	An object O of a type with extensibility kind FINAL.
O : APPENDABLE_TYPE	An object O of a type with extensibility kind APPENDABLE. This is the default for collection types and structured types.

7.4.3.5.2 Encoding of Optional Members

PLAIN_CDR serializes optional members by prepending either a ShortMemberHeader or a 12 byte LongMemberHeader. See Clause 7.4.1.1.5.2. The associated size is set to zero if the optional member is not present or to the actual serialized size if the member is present. These headers are serialized at a 4-byte offset relative to the current stream origin (XCDR.origin) and adjust the alignment origin to zero for the serialization of the member itself.

PLAIN_CDR2 and DELIMITED_CDR serialize optional members by first serializing a boolean (<is_present>) that indicates whether the member is present or not. The serialized boolean shall be set to 0 if the member is not present and to 1 if it is. If the member present (<is_present> = 1) it shall be serialized following the <is_present> boolean. If it is not present, the member shall be omitted from the serialization.

PL_CDR and PL_CDR2 serialize optional members as it would with regular members except that if the optional member is not present, then the corresponding member header and serialized member are omitted from the serialized stream.

7.4.3.5.3 Complete Serialization Rules

(1) XCDR << {O : TOP_LEVEL_TYPE} =

```
XCDR
  << INIT( OFFSET=0, ORIGIN=0,
          CENDIAN=<E>, EVERSION=<eversio> )
  << { ENC_HEADER(<E>, <eversio>, O.type) : Byte[2] }
  << PUSH( EVERSION = <eversio> )
  << PUSH( MAXALIGN = MAXALIGN(<eversio>) )
  << { <OPTIONS> : Byte[2] }
  << { O : AsNested(O.type) }
```

(2) XCDR << {O : PRIMITIVE_TYPE} =

```
XCDR
  << ALIGN( O.ssize )
  << ESWAP( AsBytes(O) )
```

```
(3) XCDR << {O : STRING_TYPE} =  
    XCDR  
    << { O.ssize : UInt32 } // includes NUL  
    << { O[i] : Byte }*    // includes NUL
```

```
(4) XCDR << {O : WSTRING_TYPE} =  
    XCDR  
    << { O.ssize : UInt32 } // No NUL  
    << { O[i] : Char16 }*  // No NUL
```

```
(5) XCDR << {O : ENUM_TYPE} =  
    XCDR  
    << { O.value : O.holder_type }
```

```
(6) XCDR << {O : BITMASK_TYPE} =  
    XCDR  
    << { O.value : O.holder_type }
```

```
(7) XCDR << {O : ALIAS_TYPE} =  
    XCDR  
    << { O : O.base_type }
```

// Arrays of primitive element type (version 1 and 2 encoding)

```
(8) XCDR << {O : PARRAY_TYPE} =  
    XCDR  
    << { O[i] : O.element_type }*
```

// Arrays (any extensibility) using version 2 encoding

(9) XCDR[2] << {O : ARRAY_TYPE} =

XCDR

<< { DHEADER(O, <E>) : UINT32 }

<< PUSH (CENDIAN = <E>)

<< { O[i] : O.element_type }*

// Arrays (any extensibility) using version 1 encoding

(10) XCDR[1] << {O : ARRAY_TYPE} =

XCDR

<< { O[i] : O.element_type }*

// Arrays with extensibility APPENDABLE use common APPENDABLE rules:

// (29)-(30)

// Arrays with extensibility MUTABLE are not allowed. Treated as APPENDABLE.

// Sequences of primitive element type (version 1 and 2 encoding)

(11) XCDR << { O : PSEQUENCE_TYPE } =

XCDR

<< { O.length : UInt32 }

<< { O[i] : O.element_type }*

// Sequences (any extensibility) using version 2 encoding

(12) XCDR[2] << {O : SEQUENCE_TYPE} =

XCDR

<< { DHEADER(O, <E>) : UINT32 }

<< PUSH (CENDIAN = <E>)

<< { O.length : UINT32 }

<< { O[i] : O.element_type }*

// Sequences (any extensibility) using version 1 encoding

```
(13) XCDR[1] << {O : SEQUENCE_TYPE} =  
    XCDR  
        << { O.length : UInt32 }  
        << { O[i] : O.element_type }*
```

// Sequences with extensibility APPENDABLE use common APPENDABLE rules:

// (29)-(30)

// Sequences with extensibility MUTABLE are not allowed. Treated as
// APPENDABLE.

// Maps of primitive key and element type (version 1 and 2 encoding)

```
(14) XCDR << {O : PMAP_TYPE} =  
    XCDR  
        << { O.length : UInt32 }  
        << { (O[i].key : O.key_type),  
            (O[i].element : O.element_type) }*
```

// Maps (any extensibility) using version 2 encoding

```
(15) XCDR[2] << { O : MAP_TYPE } =  
    XCDR  
        << { DHEADER(O, <E>) : UINT32 }  
        << { O.length : UINT32 }  
        << { (O[i].key : O.key_type),  
            (O[i].element : O.element_type) }*  
        << POP (CENDIAN)
```



```

// Maps (any extensibility) using version 1 encoding
(16) XCDR[1] << {O : MAP_TYPE} =
    XCDR
        << { O.length : UInt32 }
        << { (O[i].key : O.key_type),
            (O[i].element : O.element_type) }*

// Maps with extensibility APPENDABLE use common APPENDABLE rules:
// (29)-(30)
// Maps with extensibility MUTABLE are not allowed. Treated as APPENDABLE.

// Structures with extensibility FINAL (version 1 and 2 encoding)
// FMEMBER can be NOPT_FMEMBER (18) or OPT_FMEMBER (19)
(17) XCDR << {O : FSTRUCT_TYPE} =
    XCDR
        << { O.member[i] : FMEMBER }*

// Non-optional member of final Aggregated type (structure, union)
(18) XCDR << {M : NOPT_FMEMBER} =
    XCDR
        << { M.value : M.value.type }

// Optional member of final Aggregated type (structure, union), version 1
// see (26) and (27) for MMEMBER serialization
(19) XCDR[1] << {M : OPT_FMEMBER} =
    XCDR
        << { M : MMEMBER }

```

// Optional member of final aggregated type (structure, union), version 2

(20) XCDR[2] << {M : OPT_FMEMBER} =

XCDR

<< { <is_present> : BOOLEAN }

<< IF (<is_present>) { M.value : M.value.type }

// Structures extensibility APPENDABLE handled by generic APPENDABLE rules:

// (29)-(30)

// Structures with extensibility MUTABLE, version 2 encoding

(21) XCDR[2] << {O : MSTRUCT_TYPE} =

XCDR

<< { DHEADER(O, <E>) : UInt32 }

<< PUSH (CENDIAN = <E>)

<< { O.member[i] : MMEMBER }*

<< POP (CENDIAN)

// Member of mutable aggregated type (structure, union), version 2 encoding

(22) XCDR[2] << {M : MMEMBER} =

XCDR

<< { EMHEADER1(M) : UInt32 }

<< IF (LC(M)>=4) { NEXTINT(M) : UInt32 }

<< IF (LC(M)>=5) XCDR.offset = XCDR.offset-4

<< { M.value : M.value.type }

// Structures with extensibility MUTABLE, version 1 encoding

(23) XCDR[1] << {O : MSTRUCT_TYPE} =

```
XCDR  
<< { O.member[i] : MMEMBER }*  
<< { PID_SENTINEL : UInt16 }  
<< { length = 0 : UInt16 }
```

// Member of mutable aggregated type (structure, union), version 1 encoding

// using short PL encoding when both M.id <= 2^14 and M.value.ssize <= 2^16

(24) XCDR[1] << {M : MMEMBER} =

```
XCDR  
<< ALIGN(4)  
<< { FLAG_I + FLAG_M + M.id : UInt16 }  
<< { M.value.ssize : UInt16 }  
<< PUSH( ORIGIN=0 )  
<< { M.value : M.value.type }
```

// Member of mutable aggregated type (structure, union), version 1 encoding

// using long PL encoding

(25) XCDR[1] << {M : MMEMBER} =

```
XCDR  
<< ALIGN(4)  
<< { FLAG_I + FLAG_M + PID_EXTENDED : UInt16 }  
<< { slength=8 : UInt16 }  
<< { M.id : UInt32 }  
<< { M.value.ssize : UInt32 }  
<< PUSH( ORIGIN=0 )  
<< { M.value : M.value.type }
```

```
// Unions with extensibility FINAL (version 1 and 2 encoding)  
// see (18) to (20) for NOPT_FMEMBER and FMEMBER serialization
```

```
(26) XCDR << {O : FUNION_TYPE} =  
    XCDR  
        << { O.disc : NOPT_FMEMBER }  
        << { O.selected_member : FMEMBER }?
```

```
// Unions extensibility APPENDABLE handled by generic APPENDABLE rules:  
// (29)-(30)
```

```
// Unions with extensibility MUTABLE, version 2 encoding  
// see (22) for serialization of MMEMBER using version 2 encoding
```

```
(27) XCDR[2] << {O : MUNION_TYPE} =  
    XCDR  
        << { DHEADER(O, <E>) : UInt32 }  
        << PUSH (CENDIAN = <E>)  
        << { O.disc : MMEMBER }  
        << { O.selected_member : MMEMBER }?  
        << POP (CENDIAN)
```

```
// Unions with extensibility MUTABLE, version 1 encoding  
// see (25)-(26) for serialization of MMEMBER using version 1 encoding
```

```
(28) XCDR[1] << {O : MUNION_TYPE} =  
    XCDR  
        << { O.disc : MMEMBER }  
        << { O.selected_member : MMEMBER }?  
        << { PID_SENTINEL : UInt16 }  
        << { length = 0 : UInt16 }
```

// Extensibility APPENDABLE (Collection or Aggregated types), version 1

// encoding

(29) XCDR[1] << {O : APPENDABLE_TYPE} =

XCDR

<< { O : AsFinal(O.type) }

// Extensibility APPENDABLE (Collection or Aggregated types), version 2

// encoding

(30) XCDR[2] << {O : APPENDABLE_TYPE} =

XCDR

<< { DHEADER(O, <E>) : UInt32 }

<< PUSH (CENDIAN = <E>)

<< { O : AsFinal(O.type) }

<< POP (CENDIAN)

7.4.4 XML Data Representation

The XML Data Representation provides for the serialization of individual data samples in XML.

Each data sample shall constitute a separate XML document. The structure of that document shall conform to the XML Schema Type Representation for the sample's corresponding type definition.

(Note that, unlike in the CDR Representation, samples of mutable types are serialized no differently than samples of final or appendable types.)

The XML Data Representation has two variants: the Valid XML Data Representation and the Well-formed XML Data Representation. Their specifications follow. They both make use of the following non-normative example type definitions:

```
module MyModule1 { module MyModule2 {
    @nested
    struct MyInnerStructure {
        long my_integer;
    };
    struct MyStructure {
        MyInnerStructure inner;
        sequence<double> my_sequence_of_doubles;
    };
}}
```

7.4.4.1 Valid XML Data Representation

The XML document shall declare the namespace(s) against which it may be validated. In the event that the XSD Type Representation of the sample's type does not specify an explicit target namespace, the modules that scope that type shall imply the namespace for the document. This implied namespace shall take the form `ddstype://www.omg.org/<module path>`, where `<module path>` is a list of enclosing modules, separated by forward slashes, from outermost to innermost. The namespace prefix is not specified.

For example, the Valid XML Data Representation of an object of the example type defined above would be as follows:

```
<my:MyStructure xmlns:my="ddstype://www.omg.org/MyModule1/MyModule2">
  <my:inner>
    <my:my_integer>5<my:my_integer>
  </my:inner>
  <my:my_sequence_of_doubles>
    <my:item>10.0</my:item>
    <my:item>20.0</my:item>
    <my:item>30.0</my:item>
  </my:my_sequence_of_doubles>
</my:MyStructure>
```

7.4.4.2 Well-formed XML Data Representation

The XML document shall *not* declare the namespace(s) against which it may be validated, regardless of whether a target namespace was specified in the XSD Type Representation of the corresponding sample's type. In other words, the document shall be well-formed but not valid. This limitation allows the document to be more compact in cases where the namespace is not needed or can be inferred by the recipient.

For example, the Well-formed XML Data Representation of an object of the example type defined above would be as follows:

```
<MyStructure>
  <inner>
    <my_integer>5<my_integer>
  </inner>
  <my_sequence_of_doubles>
    <item>10.0</item>
    <item>20.0</item>
    <item>30.0</item>
  </my_sequence_of_doubles>
</MyStructure>
```

Non-normative note: Valid XML data representation can be nearly as compact as the well-formed XML data presentation by using a default namespace. The syntax to select the default namespace is `xmlns="ddstype://www.omg.org/..."`. No prefix is necessary at every element name as they now default to the default namespace. For really small datatypes (e.g., a 2d point) even the overhead of including the default namespace may be non-trivial. In such cases, well-formed XML data presentation may be preferred.

7.5 Language Binding

The Language Binding Module specifies the alternative programming-language mechanisms an application can use to construct and introspect types as well as objects of those types. These mechanisms include a Dynamic API that allows an application to interact with types and data without compile-time knowledge of the type. Note that language-specific PSMs might overrule some or all of the language binding rules specified below.

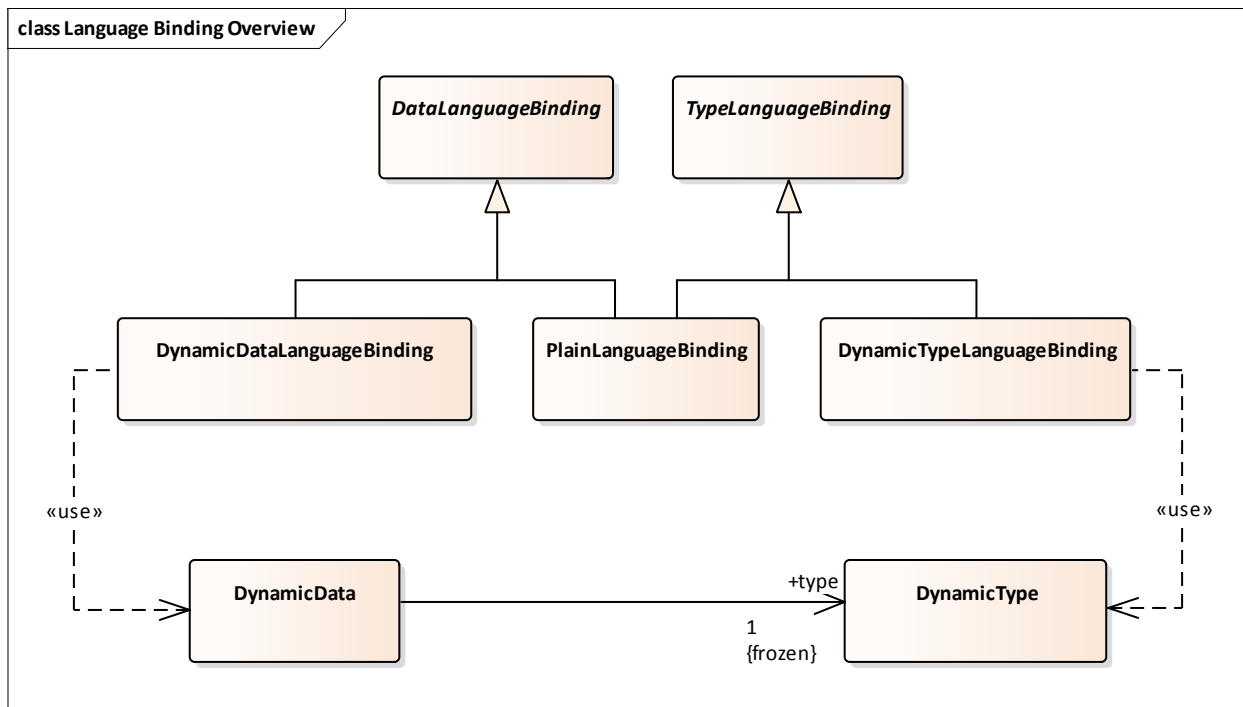


Figure 25 – Language Bindings—conceptual model

The specification defines two language bindings: Plain Language Objects and Dynamic Data. The main characteristics and motivation for each of these bindings are described in Table 41.

The Type Language Binding provides an API to manipulate types. This includes constructing new types as well as introspecting existing types. The API is the same regardless of the type, allowing applications to manipulate types that were not known at compile time. This API is similar in purpose to the `java.lang.Class` class in Java.

The principal mechanism to interact with a Type is the `DynamicType` interface. This interface is described in Clause 7.5.

Table 41 – Kinds of Language Bindings

<i>Data Representation</i>	<i>Description</i>	<i>Reasons and drawbacks</i>
Plain Language Binding	<p>Each data type is mapped into the most natural “native” construct in the programming language of choice.</p> <p>For example a STRUCT type is mapped into a class in Java where each member of the STRUCT appears as a field in the class.</p>	<p>Advantages:</p> <ul style="list-style-type: none"> • Natural - Well integrated in the programming language • Very compact notation • Very efficient <p>Disadvantages</p> <ul style="list-style-type: none"> • Requires compile-time knowledge of the data type • Changes require recompilation • Support for type evolution and sparse data can be cumbersome
Dynamic Language Binding	<p>All data types are mapped into a single Language “Dynamic Data” construct which contains operations to do introspection and access the data within.</p>	<p>Advantages:</p> <ul style="list-style-type: none"> • Does not require compile-time knowledge of the data type • Does not require code-generation • Well suited for type evolution and sparse data <p>Disadvantages</p> <ul style="list-style-type: none"> • No compile-time checking • More cumbersome to use than plain data objects • May be lower performance to use than plain data objects

7.5.1 Plain Language Binding

This mapping reuses the OMG-standard IDL language mappings [C-MAP, C++-MAP, JAVA-MAP]. It extends the most commonly used of these bindings in order to express the extended IDL constructs defined in this specification.

The following steps define this language binding in all supported programming language for a particular type.

1. First, express the type in IDL as specified in Clause 7.3.1.
2. Then, apply the OMG Standard IDL to Language Mapping to the IDL in step 2.
3. Finally, apply any programming language-specific transformations to the generated code, if applicable. These transformations are defined below.

Note that any of the following language bindings may be overridden in a language-specific PSM, such as [DDS-PSM-CXX].

7.5.1.1 Primitive Types

7.5.1.1.1 C

The Service shall provide `typedefs` with the following names to types available on the underlying platform that have the appropriate sizes and representations.

Programmers concerned with DDS portability should use the Plain Language Binding types in Table 42 below. However, some may feel that using these types impairs readability. Therefore, compliant implementations have the following degrees of freedom:

- On platforms where a native C type (e.g. `int`) is guaranteed to be identical to a DDS type, the implementation may generate the equivalent native C type.
- On platforms compliant with the C99 specification, the implementation may generate equivalent C99-compatible types.

These degrees of freedom are not expected to impact code portability, as all of these `typedefs` will map to the same underlying native C types.

Table 42 – Plain Language Binding for Primitive Types in C

<i>DDS Type</i>	<i>Plain Language Binding Type</i>	<i>Equivalent C99 Type</i>
Int32	DDS_Int32	<code>int32_t</code>
UInt32	DDS_UInt32	<code>uint32_t</code>
Int16	DDS_Int16	<code>int16_t</code>
UInt16	DDS_UInt16	<code>uint16_t</code>
Int64	DDS_Int64	<code>int64_t</code>
UInt64	DDS_UInt64	<code>uint64_t</code>
Float32	DDS_Float32	<i>(unspecified)</i>
Float64	DDS_Float64	<i>(unspecified)</i>
Float128	DDS_Float128	<i>(unspecified)</i>
Char8	DDS_Char8	<i>(unspecified)</i>
Char16	DDS_Char16	<i>(unspecified)</i>

Boolean	DDS_Boolean	Bool
Byte	DDS_Byte	<i>(unspecified)</i>

With respect to `DDS::Boolean`, only the values 0 and 1 are defined. Other values result in unspecified behavior.

With respect to `DDS::Char16`, compliant implementations may consider `wchar_t` to be an equivalent C type if the platform supports it and it is of sufficient size. Otherwise, they may map `Char16` to an equivalent integer type.

7.5.1.1.2 C++

The Service shall provide `typedefs` with the following names to types available on the underlying platform that have the appropriate sizes and representations.

Programmers concerned with DDS portability should use the Plain Language Binding types in Table 43 below. However, some may feel that using these types impairs readability. Therefore, compliant implementations have the following degrees of freedom:

- On platforms where a native C++ type (e.g. `int`) is guaranteed to be identical to a DDS type, the implementation may generate the equivalent native C++ type.
- On platforms compliant with the C99 specification, the implementation may generate equivalent C99-compatible types.

Table 43 – Plain Language Binding for Primitive Types in C++

<i>DDS Type</i>	<i>Plain Language Binding Type</i>	<i>Equivalent C99 Type</i>
Int32	DDS::Int32	[std::]int32_t
UInt32	DDS::UInt32	[std::]uint32_t
Int16	DDS::Int16	[std::]int16_t
UInt16	DDS::UInt16	[std::]uint16_t
Int64	DDS::Int64	[std::]int64_t
UInt64	DDS::UInt64	[std::]uint64_t
Float32	DDS::Float32	<i>(unspecified)</i>
Float64	DDS::Float64	<i>(unspecified)</i>
Float128	DDS::Float128	<i>(unspecified)</i>
Char8	DDS::Char8	<i>(unspecified)</i>
Char16	DDS::Char16	<i>(unspecified)</i>
Boolean	DDS::Boolean	bool <i>or</i> Bool

Byte	DDS::Byte	(<i>unspecified</i>)
------	-----------	------------------------

With respect to `DDS::Boolean`, only the values 0 and 1 are defined. Alternatively, the C++ keywords `true` and `false` may be used. Other values result in unspecified behavior.

With respect to `DDS::Char16`, compliant implementations may consider `wchar_t` to be an equivalent C++ type if the platform supports it and it is of sufficient size. Otherwise, they may map `Char16` to an equivalent integer type. This means that `DDS::Char16` may not be distinguishable from integer types for purposes of overloading.

Types `DDS::Boolean`, `DDS::Char8`, and `DDS::Byte` may all map to the same underlying C++ type. This means that these types may not be distinguishable for the purposes of overloading.

All other mappings for basic types shall be distinguishable for the purposes of overloading. That is, one can safely write overloaded C++ functions for `DDS::Int16`, `DDS::UInt16`, `DDS::Int32`, and so on.

7.5.1.2 Annotations and Built-in Annotations

IDL annotations, including the built-in annotations, impact the language binding as defined below.

7.5.1.2.1 Enumerated Literal Values

Literals in an enumerated type may be given explicit values, as defined in Clause 7.2.2.4.1. This addition to the language impacts the bindings for C, C++, and Java in the following ways.

7.5.1.2.1.1 C

The OMG-standard IDL mapping to C language [C-MAP] transforms an IDL enumeration into a series of `#define` directives, each corresponding to one of the literals in the enumeration. The values to which these definitions correspond shall be the actual values of the enumerated literals on which the definitions are based, whether implicitly or explicitly defined.

7.5.1.2.1.2 C++

The OMG-standard IDL mapping to C++ mapping [C++-MAP] transforms an IDL enumeration into a C++ enumeration. The C++ programming language supports custom values for enumerated literals. Therefore, for any enumerated literal in IDL that bears the `Value` annotation, the corresponding C++ enumerated literal definition shall be followed by an equals sign (`=`) and the value of the data member of the annotation.

7.5.1.2.1.3 Java

The OMG-standard IDL mapping to Java [JAVA-MAP] uses the pre-Java 5 “type-safe enumeration” design pattern. The value of each IDL enumerated literal is given in a Java integer constant of the following form:

```
public static final int _<label> = <value>;
```

...where *<label>* is the name of the IDL constant and *<value>* is its numeric value. As per *this* specification, that numeric value shall be set according to the explicit or implicit value assigned according to the operative Type Representation.

7.5.1.2.2 Bitmask Types

The language binding for bitmask types is defined based on the language binding for enumerations.

For each bitmask type defining flags FLAG_0 through FLAG_n, the language binding shall be as if there was an enumeration definition like the following:

```
@bit_bound(<bit_bound_value>)
enum <TypeName>Bits {
    @value(1 << <flag_value_0>)
    FLAG_0,
    ...
    @value(1 << <flag_value_n>)
    FLAG_n,
};
```

Furthermore, the language binding shall be as if there was a typedef like the following, used to represent collections of flags from the previously defined enumeration:

```
typedef <unsigned_integer_equivalent> <TypeName>;
```

...where the type *<unsigned_integer_equivalent>* is chosen based on the bound of the bitmask type as defined in Table 44 below.

Table 44 – Bit mask integer equivalents

<i>Bitmask Bound</i>	<i>Unsigned Integer Equivalent</i>
1–8	octet
9–16	unsigned short
17–32	unsigned long
33–64	unsigned long long

For example, consider the following IDL definition:

```
@bit_bound(19)
bitmask MyFlags {
    FIRST_FLAG,
    @position(14)
    SECOND_FLAG,
    THIRD_FLAG,
};
```

The language binding shall be as if the previous definition were replaced by the following:

```
enum MyFlagsBits {
    @value(1 << 0)
    FIRST_FLAG,
    @value(1 << 14)
    SECOND_FLAG,
    @value(1 << 15)
    THIRD_FLAG,
};
typedef unsigned long MyFlags;
```

7.5.1.2.3 External Members

The storage for a member of an aggregated type may be declared to be external to the storage of the enclosing object of that type. This is desirable, for instance, when the memory for a member may already exist somewhere and an application wants to combine it with other members and publish it as a unit without making additional copies. Another use case is sharing the data associated with the member among members in different objects.

The language bindings for C, Traditional C++, C++ for the DDS-PSM-CXX, and Java are provided in the following sub clauses.

7.5.1.2.3.1 C

External members shall be represented using pointers. Specifically:

- String and wide string members are already represented using pointers, so the mappings for these members do not change. The same applies to aliases to string and wide string types.
- Other external members are mapped like non-external members except that a member of type *X* shall instead be mapped as type *pointer-to-X*. For example, `short` shall be replaced by `short*`.

The constructor/initializer of the enclosing object shall set the external member pointers to NULL.

The destructor of the enclosing object shall delete the objects referenced by non-NULL external member pointers. It is the responsibility of the application to set the external member pointers to NULL before destroying the enclosing object if they do not want to delete specific referenced objects.

The copy function of the enclosing object shall do a deep copy of the external members. If the destination external member is NULL it shall be allocated. If the destination external member is not NULL it shall be filled with a copy of the source member (i.e. perform logically a recursive call to `copy(destination->pointer-to-X, source->pointer-to-X)`). If the (recursive call to the) copy operation of the external member fails, then the copy function of the containing

object shall fail as well. This may happen when the destination member is not large enough to hold a copy of the source.

There may be an additional copy function that takes in arguments which allow the user to control the behavior of the copy operation. This additional copy function shall allow the user to choose whether a shallow or deep copy is made as well as whether any existing memory pointed by the member is reused, released, or replaced during the copy.

In the case that a shallow copy is made and the destination member is NULL then the destination member pointer will be set to the source member pointer.

In the case that a deep copy is made and the destination member pointer is NULL, memory for the destination member will be allocated and then copied into.

For the behaviors supported by the additional copy function when the destination member is not NULL, see Table 45.

Table 45 – Configurable behaviors of the copy function when destination is not NULL

<i>Copy Type</i>	<i>Action when destination member is not NULL</i>	<i>Description</i>
Shallow Copy	Replace	Destination will now point to the same memory address as source. The existing memory pointed to by destination is released before making the assignment.
	Release	Destination will now point to the same memory address as source. The existing memory pointed to by destination is released before making the assignment.

Deep Copy	Reuse	(Default) Try to reuse the existing memory to copy into. If the existing member is not large enough, this operation shall fail.
	Replace	Replace the destination member. Allocate new memory to copy into and replace the existing memory without releasing it. It is the application's responsibility to release the replaced memory.
	Release	Release the existing memory before allocating new memory to copy into.

7.5.1.2.3.1.1 External Optional Members

A member that is both external and optional shall be mapped as if it was just external. The difference is that it is valid for the member to be NULL when writing a sample containing this member. If the member is only external but not optional, then it is not allowed for the member to be NULL at the time of a write.

7.5.1.2.3.2 Traditional C++

This mapping extends the IDL to C++ language mapping defined in [C++-MAP].

External members shall be represented by any type that behaves similarly to a pointer (e.g., a plain pointer or a `_var` type). The chosen type must support the concept of being “unset.” For example, a plain pointer is considered unset if its value is NULL.

- In cases where the non-external mapping already uses a type similar to a pointer, it shall remain unchanged.
- In cases where the non-external mapping uses a member of type `x`, `x` shall be replaced by `pointer-to-x`. For example, if plain pointers are used, `short` shall be replaced by `short*`.

The behavior of the constructor, destructor, and copy functions shall be the same as specified for C.

7.5.1.2.3.2.1 External Optional Members

A member that is both external and optional shall be mapped as if it was just external. The difference is that it is valid for the member to be unset when writing a sample containing this member. If the member is only external but not optional, then it is not allowed for the member to be unset at the time of a write.

7.5.1.2.3.3 Modern C++

This mapping extends the IDL to C++ language mapping defined in [DDS-PSM-CXX].

External members shall be represented as an instantiation of a template class `external<T>`, where `T` is the type of the external member. This is a “smart pointer” class that wraps a shared pointer, `ptr_` for automatic reference counting and a boolean `locked_` that controls the assignment behavior. The destruction of the object referenced by an external member is always managed by the underlying shared pointer.

The value of the `locked_` attribute dictates whether copying an external member performs a deep copy or shallow copy of the referenced member. It can also be used to prevent sharing of the referenced object. This control is useful in some situations, for example, to prevent sharing a reference to memory that belongs to a `DataReader` in a DDS application. See Sub Clauses 7.5.1.2.3.3.4 and 7.5.1.2.3.3.5 for details about the copy constructor and assignment operator.

The `locked_` attribute is set at the time the external member is constructed and cannot be modified. The `locked_` attribute can only be set to `true` when the shared pointer is set to a non-NULL value.

The `external<T>` class shall be generated inside of an appropriate namespace. In the case of [DDS-PSM-CXX], this namespace is `dds::core`.

```
namespace dds { namespace core {
template <typename T>
class external {
public:
    external();
    external(T* p, bool locked = false);
    external(shared_ptr<T> p);
    external(const external& other);
    ~external();
    external& operator=(const external& other);
    T& operator*();
    const T& operator*() const;
    T* get();
    const T* get() const;
    shared_ptr<T> get_shared_ptr();
    T* operator->();
    const T* operator->() const;
    bool operator==(const external<T>& other) const;
    bool operator!=(const external<T>& other) const;
    operator bool() const;
    bool is_locked() const;
};
};
};
```



```

    void lock();

private:
    shared_ptr<T> ptr_;
    bool locked_;
};
}} // namespace dds::core

```

7.5.1.2.3.3.1 Operation: Default Constructor

Create an empty `external<T>` object with an empty `ptr_` and `locked_` initialized to false.

7.5.1.2.3.3.2 Operation: Constructor from a T*

Create a new `external<T>` object referencing the provided managed object. The attribute `locked_` is set to false and `ptr_` is initialized with `p`.

Parameter `p` - The object for `ptr_` to manage.

Parameter `locked` - Whether or not the constructed `external<T>` should be locked. This is an optional parameter with a default value of false.

7.5.1.2.3.3.3 Operation: Constructor from a shared pointer to T object

Create a new `external<T>` object that references the same object managed by the specified shared pointer `p`. The attribute `locked_` is set to false and `ptr_` is initialized with `p`.

Parameter `p` - The `shared_ptr<T>` holding the `T*` reference that will be shared with the new `external<T>` object.

7.5.1.2.3.3.4 Operation: Copy Constructor

Creates an external object from an existing external object (`other`). The behavior of this operation depends on the value of the `locked_` attribute of the existing external object (`other`).

- If `other.is_locked()` is false, then the new `external<T>` object shares the reference with `other`. In other words this operation will not create a `T` object, instead it will perform a shallow copy of `T*` pointer.
- If `other.is_locked()` is true, then a new `T` object is created and `ptr_` is initialized with a reference to the newly created `T` object. The contents of newly-allocated object are initialized with a copy from the contents of `other`. In other words this operation will create a new `T` object and do a deep copy.

Either way, the newly constructed `external<T>` object will have `locked_` set to false.

Parameter `other` - The external object used to initialize the new constructed `external<T>` object.

7.5.1.2.3.3.5 Operation: Assignment Operator

Assigns an external object to another.

The behavior of this operation depends on the value of the `locked_` attribute both on the source of the copy as well as on the destination.

The behavior specified in Table 46 below shall be applied when assigning an `external<T>` object source to another `external<T>` object destination:

Table 46 – Behavior of assignment operator

<i>Destination</i>	<i>Destination</i>	<i>Source</i>	<i>Source</i>	<i>Behavior of assignment operator</i>
<code>locked_</code>	<code>ptr_</code>	<code>locked_</code>	<code>ptr_</code>	
TRUE	<any>	<any>	<any>	Error. Operation cannot be called when <code>destination.is_locked() == TRUE</code>
FALSE	<any>	<any>	EMPTY	The destination is reset. Result is <code>destination.ptr_</code> is EMPTY.
FALSE	EMPTY	TRUE	Not EMPTY	Create new object for <code>destination.ptr_</code> Perform deep copy from <code>source.ptr_</code> to <code>destination.ptr_</code> .
FALSE	Not EMPTY	TRUE	Not EMPTY	Reuse existing <code>destination.ptr_</code> Perform deep copy from <code>source.ptr_</code> into the existing <code>destination.ptr_</code> .
FALSE	<any>	FALSE	Not EMPTY	Perform shallow copy. The <code>destination.ptr_ == source.ptr_</code> Destination will reference same object as source

Parameter `other` - The external object whose contents are assigned to this external object.

7.5.1.2.3.3.6 Operation: Destructor

Destroy the external object. If `ptr_` is the last reference to the managed object, then the managed object will be released, otherwise the reference count will simply be decreased.

7.5.1.2.3.3.7 Operation: `operator*` (const and non-const versions)

Get a reference to the underlying managed object that `ptr_` points at.

7.5.1.2.3.3.8 Operation: `get` (const and non-const versions)

Obtains a pointer to the managed object.

7.5.1.2.3.3.9 Operation: `get_shared_ptr`

Obtains a shared pointer to the managed object.

7.5.1.2.3.3.10 Operation: operator-> (const and non-const versions)

Allows accessing members of the managed object.

7.5.1.2.3.3.11 Operation: operator==

Returns whether two external objects manage the same object or are both empty.

7.5.1.2.3.3.12 Operation: operator!=

Returns whether two external objects do not manage the same object.

7.5.1.2.3.3.13 Operation: operator bool

Checks if there is a managed object (is not NULL) or not (is NULL).

7.5.1.2.3.3.14 Operation: is_locked

Indicates whether this object is locked or not.

7.5.1.2.3.3.15 Operation: lock

Sets the `locked_` attribute to true. This prevents of the `external<T>` object from modifying the referenced `T` object. This means that future assignment operations to the `external<T>` object will fail and any copies from `external<T>` will be deep copies (i.e., not share a reference to the same underlying `T` object).

7.5.1.2.3.3.16 External Optional Members

A member that is both external and optional shall be mapped as if it was just external. The difference is that it is valid for `ptr_` to be empty when writing a sample containing this member. If the member is only external but not optional, then it is not allowed for `ptr_` to be empty at the time of a write.

7.5.1.2.3.4 Java

This mapping extends the IDL to Java language mapping defined in [JAVA-MAP].

External members shall be represented using object references. Since *all* objects are referred to by reference in Java, the mappings for external members of non-primitive types are identical to those of non-external members. For IDL types that map to Java primitive types, those Java primitive types shall be replaced by the corresponding object box types from the `java.lang` package. For example, `short` shall be replaced by `java.lang.Short`.

7.5.1.2.4 Optional Members

A member of an aggregated type may be declared to be optional, meaning that its value may be omitted from sample to sample of that type. This concept impacts the language bindings for C, C++, and Java in the following ways.

7.5.1.2.4.1 C

Optional members shall be represented using pointers. Specifically:

- String and wide string members are already represented using pointers, so the mappings for these members shall not change. The same shall apply to aliases to string and wide string types.
- Other optional members are mapped like non-optional members except that a member of type *X* shall instead be mapped as type *pointer-to-X*. For example, `short` shall be replaced by `short*`.

A `NULL` pointer shall indicate an omitted value.

7.5.1.2.4.2 C++

Optional members shall be represented using plain pointers rather than automatic values or smart pointers.

- In cases where the mapping of non-optional members already uses a plain pointer, it shall remain unchanged.
- In cases where the mapping of non-optional members uses a “`_var`” smart pointer, the `_var` type shall be replaced by the corresponding plain pointer type. For example, `MyType_var` is replaced by `MyType*`.
- In cases where the mapping of non-optional members uses an automatic member of type *X*, *X* shall be replaced by *pointer-to-X*. For example, `short` shall be replaced by `short*`.

A `NULL` pointer shall indicate an omitted value.

7.5.1.2.4.3 Java

Optional members shall be represented using object references. Since *all* objects are referred to by reference in Java, the mappings for optional members of non-primitive types are identical to those of non-optional members. For IDL types that map to Java primitive types, those Java primitive types shall be replaced by the corresponding object box types. For example, `short` shall be replaced by `java.lang.Short`.

A `null` pointer shall indicate an omitted value.

7.5.1.2.4.4 Optional Arrays in C and C++

Optional arrays having element type “*T*” shall be mapped to type *pointer-to-array-of-type-T* rather than to type *array-of-pointers-to-type-T*.

For example, the structure `MyStruct` containing an optional array of ten integers defined by the IDL:

```
// IDL declaration
struct MyStruct {
    @optional long array_member[10];
};
```

Should be mapped in C and C++ to the type:

```
// Mapping to C/C++
struct MyStruct {
    int32_t (*array_member)[10];
}
```

Without the parentheses, `array_member` is an array of ten `int32_t` pointers, rather than a pointer to an array of ten `int32_t` values.

7.5.1.2.5 Nested Types

An IDL compiler need not (although it may) generate `TypeSupport`, `DataReader`, or `DataWriter` classes for any nested type.

7.5.1.2.6 User-Defined Annotation Types

A type designer may define his or her own annotation types. The language bindings for these shall be as follows in Java. In programming languages that lack the concept of annotations, an implementation of this specification may choose to ignore user-defined annotations with respect to this language binding.

7.5.1.2.6.1 Java

Each user-defined IDL annotation type shall be represented by a corresponding Java annotation type. An IDL annotation type defining operations `op_1` through `op_n` shall be represented by the following Java annotation types:

```
public @interface <TypeName> {
    <op_1_type> <op_1_name>() [default <default>];
    ...
    <op_n_type> <op_n_name>() [default <default>];
}
```

```
public @interface <TypeName>Group {
    <TypeName>[] value();
}
```

The `<op_type>` shall be the Java type corresponding to the return type of the IDL operation. If a default value is specified for a given member, it shall be reflected in the Java definition. Otherwise, the Java definition shall have no default value.

A Java annotation type may itself be annotated (for example, by annotation types in the `java.lang.annotation` package). The presence or absence of any such annotations is undefined.

For each IDL element to which a single instance user-defined annotation is applied, the corresponding Java element shall be annotated with the Java annotation of the same name. For each IDL element to which multiple instances of the annotation are applied, the corresponding

Java element shall be annotated with the generated annotation bearing the “Group” suffix; each application of the user-defined annotation shall correspond to a member of the array in the group.

7.5.1.3 Map Types

The language bindings for C, Traditional C++, C++ for the DDS-PSM-CXX, and Java are provided in the following sub clauses.

Implementations are only required to support keys of types `UInt32`, `UInt64`, and `String<Char8>`. Implementations may choose to support other key types; however, to reduce complexity, maps declared to use any other key type may not be declared as an anonymous type in the IDL. If a Type Representation compiler encounters an anonymous map with key type that it does not support, it shall fail with an error.

7.5.1.3.1 Operations

Map types support operations to create, delete, and manipulate their contents. These operations are described in Table 47 below. Each of the language bindings support logically equivalent operations which are further described below if they are not supported natively by the language.

Table 47 – Operations for `map<KeyType, ElementType>`

<i>map<KeyType, ElementType></i>		
Operations		
new		<code>map<KeyType, ElementType></code>
delete		<code>void</code>
initialize		<code>void</code>
finalize		<code>void</code>
copy		<code>ReturnCode_t</code>
	source	<code>map<KeyType, ElementType></code>
	autogrow	<code>Boolean</code>
get size		<code>unsigned int</code>
get max size		<code>unsigned int</code>
set max size		<code>ReturnCode_t</code>
	max_size	<code>unsigned int</code>
clear		<code>void</code>
insert		<code>ReturnCode_t</code>
	key	<code>KeyType</code>
	element	<code>ElementType</code>
insert or assign		<code>ReturnCode_t</code>
	key	<code>KeyType</code>
	element	<code>ElementType</code>

erase		ReturnCode t
	key	KeyType
get first		ReturnCode t
get next		ReturnCode t
	inout: entry	MapEntry
find element		ElementType
	key	KeyType
find entry		MapEntry
	key	KeyType
get pair		Boolean
	entry	MapEntry
	out: key	KeyType
	out: element	ElementType

7.5.1.3.2 C

This mapping extends the IDL to C language mapping defined in [C-MAP].

Map types shall be represented as a collection of structures that contain a member of the key type followed by a member of the element type. A set of methods which create, delete and manipulate objects of the map type shall also be generated. The name of the map type is specified in this language binding.

7.5.1.3.2.1 Map Type Name

For maps whose key type is a Primitive Type the name of the map type shall be constructed by combining the key type name with the element type name. The combination shall follow the schema below:

```
[key_type][fully_qualified_element_type]Map
```

For example, the names of the maps with element type Foo for each of the three mandatory key types would be:

```
StringFooMap
```

```
UInt32FooMap
```

```
UInt64FooMap
```

The concrete language binding is not specified, implementers may choose any language binding (e.g., a structure or a sequence) as long as its name and operations comply with what is specified here.

For any type T, the declaration and implementation of the map types having element type T and key types UInt32, UInt64, and String shall be generated alongside the implementation code for element type T.

Note: each of the following operations except for new take the map to be operated on as the first parameter.

7.5.1.3.2.2 Operation: new

Allocate a new map. If this operation fails in an implementation-specific way, this operation shall return NULL.

7.5.1.3.2.3 Operation: delete

Delete the map and all of its contents.

7.5.1.3.2.4 Operation: initialize

Initialize the map. The initial size and capacity of the map shall be 0.

7.5.1.3.2.5 Operation: finalize

Finalize the map. The entries in the map will be deleted, and both the size and maximum size set to 0.

This is equivalent to calling `clear()` followed by `set_max(0)`.

7.5.1.3.2.6 Operation: copy

Overwrite the contents of this (destination) map with the contents of another (source) map. Any entries that are not present in the source map are erased from the destination map. The source map shall not be modified by this operation.

If the size of the source map is greater than the maximum size of the destination map, the behavior depends on the `autogrow` parameter. If `autogrow` is `TRUE`, the operation shall grow the maximum size of the destination map as needed. If `autogrow` is `FALSE`, the operation shall fail and return `DDS_RETCODE_PRECONDITION_NOT_MET`. In this case the destination map shall remain unchanged.

If the size of the source is less than the maximum size of the destination, then it is left to the implementation to decide whether the maximum size of the destination map is trimmed to match the source or left unchanged.

If this operation fails in an implementation-specific way, the operation shall return

`DDS_RETCODE_ERROR`.

Parameter `source` – The map whose contents are to be copied. If this argument is `NULL`, the operation shall fail with `DDS_RETCODE_BAD_PARAMETER`.

Parameter `autogrow` – Controls the behavior in case the destination map `max_size` is insufficient to hold the source map.

7.5.1.3.2.7 Operation: get_size

Get the current size of the map. The size of the map is how many entries are currently present in the map.

7.5.1.3.2.8 Operation: `get_max_size`

Get the current maximum size of the map. The maximum size limits the number of entries the map may contain.

7.5.1.3.2.9 Operation: `set_max_size`

Set the maximum size of the map.

This operation shall fail with `DDS_RETCODE_ERROR` if it fails for any implementation-specific reason.

Parameter `max_size` – The new maximum size of the map. If the new `max_size` is less than the current size of the map, the operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`.

7.5.1.3.2.10 Operation: `clear`

Clear all of the entries from the map. The size of the map is set to 0 and the maximum size does not change.

7.5.1.3.2.11 Operation: `insert`

Insert a new entry into the map with the given key and element values. If the key already exists in the map, the operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. If successful, the size shall be increased by 1. If inserting a new entry into the map would increase the size past the current maximum size, then this operation shall fail with `DDS_RETCODE_PRECONDITION_NOT_MET`.

This operation shall fail with `DDS_RETCODE_ERROR` if it fails for any implementation-specific reason.

Parameter `key` – The key value of the entry to insert. If this argument is `NULL`, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For keys with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

Parameter `element` – The element value of the entry to insert. If this argument is `NULL`, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For elements with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

7.5.1.3.2.12 Operation: `insert_or_assign`

Insert an entry into the map with the given key and element values. If the key already exists in the map, then the corresponding element shall be replaced. If the key value did not already exist in the map, then the entry shall be inserted with the same behavior specified for the `insert` operation.

This operation shall fail with `DDS_RETCODE_ERROR` if it fails for any implementation-specific reason.

Parameter `key` – The key value of the entry to insert. If this argument is `NULL`, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For keys with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

Parameter `element` – The element value of the entry to insert. If this argument is NULL, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For elements with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

7.5.1.3.2.13 Operation: `erase`

Remove the entry with the given key from the map. If successful, the size of the map shall be decreased by 1.

Parameter `key` – The key value of the entry to erase. If this argument is NULL, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For keys with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

7.5.1.3.2.14 Operation: `get_first`

Retrieves a `MapEntry` referencing the first entry in the map. The returned `MapEntry` may be a sentinel if the map is empty.

7.5.1.3.2.15 Operation: `get_next`

Advance the `MapEntry` to the next entry in the Map. If the `MapEntry` was referencing the last entry, the `MapCursor` will be advanced to a sentinel and the operation will return `FALSE`, otherwise it will return `TRUE`.

7.5.1.3.2.16 Operation: `find_element`

Retrieve the element whose key matches the specified one from the map. If the key exists, then return the element corresponding to the key, otherwise return NULL.

Parameter `key` – The key value of the element to search for. If this argument is NULL, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For keys with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

7.5.1.3.2.17 Operation: `find_entry`

Retrieve the `MapEntry` whose key matches the specified one from the map. If the key exists, then return a `MapEntry` referencing the entry (key and element), otherwise return a sentinel.

Parameter `key` – The key value of the element to search for. If this argument is NULL, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`. For keys with primitive types, this argument shall be generated as the type and not as a pointer to the primitive type.

7.5.1.3.2.18 Operation: `get_pair`

Retrieve the key and element associated with the `MapEntry`. If the `MapEntry` was a sentinel the operation will return `FALSE`, otherwise it will return `TRUE` and fill the output parameters with references to the key and element.

Parameter `entry` – The `MapEntry` whose key and element we wish to retrieve. If this argument is NULL, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`.

Parameter `key` (output) – The key value associated with the `MapEntry`. If this argument is NULL, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`.

Parameter `element` – The element value associated with the `MapEntry`. If this argument is `NULL`, this operation shall fail and return `DDS_RETCODE_BAD_PARAMETER`.

7.5.1.3.2.19 Example (Non-Normative)

For a struct `MyElementType` defined by the IDL:

```
// IDL definition
module MyModule {
    struct MyElementType {
        // ...members
    };
};
```

The following structures and operations should be generated for `map<unsigned long, MyElementType>`:

```
struct UInt32MyModule_MyElementTypeMapElement {
    uint32_t key;
    MyModule_MyElementType element;
};

typedef sequence<UInt32MyModule_MyElementTypeMapElement>
UInt32MyModule_MyElementTypeMap;

// Operations on UInt32MyModule_MyElementTypeMap
UInt32MyModule_MyElementTypeMap* UInt32BarMap_new();
void UInt32MyModule_MyElementTypeMap_delete(
    UInt32MyModule_MyElementTypeMap *map);
void UInt32MyModule_MyElementTypeMap_initialize(
    UInt32MyModule_MyElementTypeMap *map);
void UInt32MyModule_MyElementTypeMap_finalize(
    UInt32MyModule_MyElementTypeMap *map);
DDS_ReturnCode_t UInt32MyModule_MyElementTypeMap_copy(
    UInt32MyModule_MyElementTypeMap *map,
    UInt32MyModule_MyElementTypeMap *other,
    bool autogrow);
uint32_t UInt32MyModule_MyElementTypeMap_get_size(
    UInt32MyModule_MyElementTypeMap *map);
DDS_ReturnCode_t UInt32MyModule_MyElementTypeMap_set_size(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t size);
```

```

uint32_t UInt32MyModule_MyElementTypeMap_get_max_size();
DDS_ReturnCode_t UInt32MyModule_MyElementTypeMap_set_max_size(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t max_size);
void UInt32MyModule_MyElementTypeMap_clear();
DDS_ReturnCode_t UInt32MyModule_MyElementTypeMap_insert(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t key,
    MyModule_MyElementType *element);
DDS_ReturnCode_t UInt32MyModule_MyElementTypeMap_insert_or_assign(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t key,
    MyModule_MyElementType *element);
DDS_ReturnCode_t UInt32MyModule_MyElementTypeMap_erase(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t key);
MapEntry UInt32MyModule_MyElementTypeMap_get_first(
    UInt32MyModule_MyElementTypeMap *map);
bool UInt32MyModule_MyElementTypeMap_get_next(
    UInt32MyModule_MyElementTypeMap *map,
    MapEntry *entry);
MyElementType* UInt32MyModule_MyElementTypeMap_find_element(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t key );
MapEntry UInt32MyModule_MyElementTypeMap_find_entry(
    UInt32MyModule_MyElementTypeMap *map,
    uint32_t key );
bool UInt32MyModule_MyElementTypeMap_get_pair(
    UInt32MyModule_MyElementTypeMap *map,
    MapEntry *entry,
    uint32_t *key,
    MyElementType **element);

```

7.5.1.3.3 Traditional C++

This mapping extends the IDL to C++ language mapping defined in [C++-MAP].

This C++ language binding differs only slightly from the C language binding. Instead of a C structure with accompanying functions, C++ defines a class with methods.

7.5.1.3.3.1 Map Class Name and operations

The map class shall be named the same as the C structure, see Sub Clause 7.5.1.3.2, except that it is placed in the same namespace as the element type declaration.

For example, the XTYPES map with key of type `UInt32` and element type `MyElementType` belonging to module `MyModule` would be bound to the class:

```
namespace MyModule {
class UInt32MyElementTypeMap {
public:
    UInt32MyElementTypeMap();
    ~UInt32MyElementTypeMap();
    ReturnCode_t copy(
        const UInt32MyElementTypeMap &other,
        bool autogrow = true);
    uint32_t get_size() const;
    ReturnCode_t set_size(uint32_t size);
    uint32_t get_max_size() const;
    ReturnCode_t set_max_size(uint32_t max_size);
    void clear();
    ReturnCode_t insert(
        uint32_t key,
        const MyElementType &element,
        bool replace = true );
    ReturnCode_t erase( uint32_t key );
    MapEntry get_first();
    bool get_next(MapEntry &entry);
    MyElementType* find_element( uint32_t key);
    MapEntry find_entry( uint32_t key );
    bool get_pair(
        const MapEntry &entry,
        uint32_t *key,
        MyElementType **element);
};
}
```

Refer to the C language binding for the behavior of each of the above methods, with the exceptions described below.

The C++ operation `insert` behaves as the C `insert()` if the `replace()` parameter is false and it behaves as the C `insert_or_assign()` if `replace` parameter is true.

7.5.1.3.4 Modern C++

This mapping extends the IDL to C++ language mapping defined in [DDS-PSM-CXX].

The Map type shall be bound to an instantiation of the `std::map` template. The C++ Standard [C++-LANG] defines the `std::map` container as follows:

```
namespace std {
    template<class Key,
            class T,
            class Compare    = less<Key>,
            class Allocator  = allocator<pair<const Key,T> >
    > class map;
}
```

The `std::map` template shall be instantiated with the `K` class parameter being the C++ type corresponding to the key type and the `T` parameter is the C++ type corresponding to the element type.

When a map has keys of a string type, the Compare function shall operate on the character contents of the strings; it shall not operate on the strings' pointer values (as `std::less` does). The instantiations for the Compare and Allocator parameters are otherwise undefined and may or may not take their default values.

For example, the XTYPES map with key of type `UInt32` and element type `MyElementType` belonging to module `MyModule` would be bound to the following template instantiation:

```
std::map<uint32_t, MyModule::MyElementType *>
```

7.5.1.3.5 Java

An IDL map type shall be represented in Java by an implementation of the standard `java.util.Map` interface. The implementation class to be used is not defined, nor is it defined whether Java 5+ generic syntax should be used. (The OMG-standard IDL mapping to Java [JAVA-MAP] predates Java 5, and implementations of it may retain compatibility with earlier versions of Java.)

The key objects for such maps shall be of the Java type corresponding to the IDL key element type. The value objects shall be of the Java type corresponding to the IDL value element type. If either of these Java types is a primitive type, then the corresponding object box type (e.g., `java.lang.Integer` for `int`) shall be used in its place.

7.5.1.3.6 Other Programming Languages

In all languages for which no language-specific mapping is specified, the language binding for map types shall be based on the equivalent IDL definition given in 7.4.1.1.4.

7.5.1.4 Structure and Union Types

The Plain Language Binding for structure and union types shall correspond to the IDL language mappings for IDL structures and unions as amended below.

7.5.1.4.1 Inheritance

A structure type that inherits from another shall be represented as follows.

7.5.1.4.1.1 C++

The C++ `struct` corresponding to the subtype shall publicly inherit from the C++ `struct` corresponding to the supertype.

7.5.1.4.1.2 Java

The Java class corresponding to the subtype shall extend the Java class corresponding to the supertype.

7.5.1.4.1.3 Other Programming Languages

The language binding shall be generated as if an instance of the base type were the first member of the subtype with the name “parent,” as in the following equivalent IDL definition:

```
struct <struct_name> {  
    <base_type_name> parent;  
    // ... other members  
};
```

7.5.2 Dynamic Language Binding

The Dynamic Type Language Binding provides an API to manipulate types. This includes constructing new types as well as introspecting existing types. The API is the same regardless of the Type, allowing applications to manipulate types that were not known at compile time. This API is similar in purpose to the `java.lang.Class` class in Java.

The Dynamic Data Language Binding provides an API to manipulate objects of any Type. This includes creating data objects, setting fields and getting fields, as well as accessing the Type associated with the data object. The API is the same regardless of the type of the object, allowing applications to manipulate data objects of types not known at compile time.

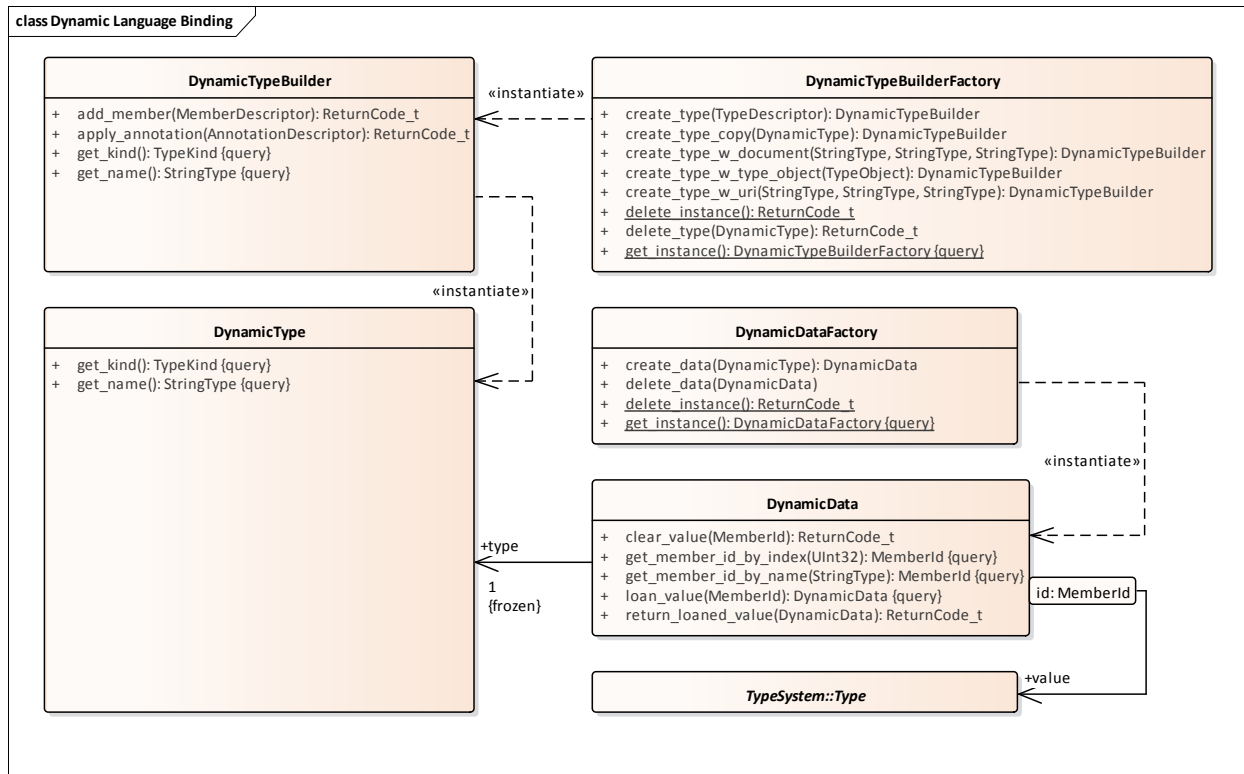


Figure 26 – Dynamic Data and Dynamic Type

There are a small number of fundamental classes to understand in this model, as well as a few helper classes:

- **DynamicType**: Objects of this class represent a type’s schema: its physical name, type kind, member definitions (if any), and so on.
- **DynamicTypeBuilderFactory**: This type is logically a singleton. Its instance is responsible for creating `DynamicType` and `DynamicTypeSupport` objects.
- **DynamicData**: A `DynamicData` object represents an individual data sample. It provides reflective getters and setters for the members of that sample.
- **DynamicDataFactory**: This type is logically a singleton. Its instance is responsible for creating `DynamicData` objects.

7.5.2.1 UML-to-IDL Mapping Rules

Each type in this Language Binding has an equivalent IDL API. These APIs are specified using the IDL Type Representation defined in this document with the addition of other standard IDL syntax. These latter parts of IDL are used to describe portions of the UML model that have requirements that go beyond those addressed by the IDL Type Representation (for example, local operations).

Specifically, UML constructs shall be mapped to IDL as described below.

- UML enumerations are mapped to IDL enumerations.

- UML classifiers with value semantics are represented as IDL valuetypes. Classifiers with reference semantics are represented as local interfaces.
- UML structural properties in most cases are represented as IDL fields or attributes.
 - Properties of classifiers mapped as valuetypes are represented as plain fields. Properties of classifiers mapped as interfaces are represented as attributes; if the property value is read-only, so is the attribute.
 - Properties with multiplicity [1] (the default if not otherwise noted) are mapped as-is.
 - Properties with multiplicity [0..1] are defined as `@optional`.
 - Properties with multiplicity [*] (equivalent to [0..*]) or [1..*] may be mapped *either* simply as sequences (in cases where the number of objects is expected to be small and the required level of abstraction low) *or*—in more complex scenarios—a set of methods:

```
unsigned long get_<property_name>_count();
DDS::ReturnCode_t get_<property_name>(
    inout <property_type> value,
    in unsigned long idx);
```

In addition, if and only if the property value can be modified:

```
DDS::ReturnCode_t set_<property_name>(
    in unsigned long idx,
    in <property_type> value);
```

The “get” operation shall fail with `RETCODE_BAD_PARAMETER` if the given index is outside of the current range. The “set” operation shall do the same with one exception: it shall allow an index one past the end (i.e., equal to the current count); setting with this index shall have the effect of appending a new value to the end of the collection. Either operation shall fail with `RETCODE_BAD_PARAMETER` if either argument is nil.

Each type mapping below indicates which of these two mappings it uses in which cases.

- Qualified association ends (representing mappings from one value to another) are mapped to a set of operations:

```
DDS::ReturnCode_t get_<property_name>(
    inout <property_type> value,
    in <qualifier_type> key);
DDS::ReturnCode_t get_all_<property_name>(
    inout map< <qualifier_type>, <property_type> > value);
```

In addition, if and only if the property value can be modified:

```

DDS::ReturnCode_t set_<property_name>(
    in <qualifier_type> key,
    in <property_type> value);

```

The “get” operation shall return with `RETCODE_NO_DATA` if no value exists for the given key. Either operation shall return with `RETCODE_BAD_PARAMETER` if either argument is nil.

- UML operations are represented as IDL operations.
 - Static operations are commented, as IDL does not formally support static operations. It is up to the implementer to reflect these operations properly in each programming language to which the IDL may be transformed.

These rules may be qualified or overridden below on a case-by-case basis. The complete IDL API can be found in “Annex C: Dynamic Language Binding.”

7.5.2.2 DynamicTypeBuilderFactory

This class is logically a singleton (although it need not technically be a singleton in practice). Its “only” instance is the starting point for creating and deleting `DynamicTypeBuilder` objects.

Table 48 – DynamicTypeBuilderFactory properties and operations

<i>DynamicTypeBuilderFactory</i>		
Operations		
static get instance		DynamicTypeBuilderFactory
static delete instance		ReturnCode t
get_primitive_type		DynamicType
	kind	TypeKind
create_type		DynamicTypeBuilder
	descriptor	TypeDescriptor
create_type_copy		DynamicTypeBuilder
	type	DynamicType
create_type_w_type_object		DynamicTypeBuilder
	type_object	TypeObject
create_string_type		DynamicTypeBuilder
	bound	UInt32
create_wstring_type		DynamicTypeBuilder
	bound	UInt32

create_sequence_type		DynamicTypeBuilder
	element_type	DynamicType
	bound	UInt32
create_array_type		DynamicTypeBuilder
	element_type	DynamicType
	bound	UInt32 [1..*]
create_map_type		DynamicTypeBuilder
	key element type	DynamicType
	element_type	DynamicType
	bound	UInt32
create_bitmask_type		DynamicTypeBuilder
	bound	UInt32
create_type_w_uri		DynamicTypeBuilder
	document_url	string<Char8>
	type_name	string<Char8>
	include_paths	string<Char8> [*]
create_type_w_document		DynamicTypeBuilder
	document	string<Char8>
	type_name	string<Char8>
	include_paths	string<Char8> [*]
delete_type		ReturnCode t
	type	DynamicType

7.5.2.2.1 Operation: create_array_type

Create and return a new `DynamicTypeBuilder` object representing an array type. All objects returned by this operation should eventually be deleted by calling `delete_type`.

All array types having equal element types, an equal number of dimensions, and equal bounds in each dimension shall be considered equal. An implementation may therefore elect whether to always return a new object from this method or whether to pool objects and to return previously created type objects consistent with these rules.

If an error occurs, this method shall return a nil value.

Parameter `element_type` – The type of all objects that can be stored in an array of the new type. If this argument is nil, the operation shall fail with `RETCODE_BAD_PARAMETER`.

Parameter bound - A collection of unsigned integers, the length of which is equal to the number of dimensions in the new array type, and the values of which are the bounds of each dimension. (For example, a three-by-two array would be described by a collection of length two, where the first element had a value of three and the second a value of two.) If this argument is nil, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.2.2 Operation: `create_bitmask_type`

Create and return a new `DynamicTypeBuilder` object representing a bitmask type. All objects returned by this operation should eventually be deleted by calling `delete_type`.

If an error occurs, this method shall return a nil value.

Parameter bound - The number of reserved bits in the bitmask. If this value is out of range, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.2.3 Operation: `create_map_type`

Create and return a new `DynamicTypeBuilder` object representing a map type. All objects returned by this operation should eventually be deleted by calling `delete_type`.

All map types having equal key and value element types and equal bounds shall be considered equal. An implementation may therefore elect whether to always return a new object from this method or whether to pool objects and to return previously created type objects consistent with these rules.

If an error occurs, this method shall return a nil value.

Parameter key_element_type - The type of all objects that can be stored as keys in a map of the new type. If this argument is nil, the operation shall fail with `RETCODE_BAD_PARAMETER`.

Parameter element_type - The type of all objects that can be stored as values in a map of the new type. If this argument is nil, the operation shall fail with `RETCODE_BAD_PARAMETER`.

Parameter bound - The maximum number of key-value pairs that may be stored in a map of the new type. If this argument is equal to `LENGTH_UNLIMITED`, the map type shall be considered to be unbounded.

7.5.2.2.4 Operation: `create_sequence_type`

Create and return a new `DynamicTypeBuilder` object representing a sequence type. All objects returned by this operation should eventually be deleted by calling `delete_type`.

All sequence types having equal element types and equal bounds shall be considered equal. An implementation may therefore elect whether to always return a new object from this method or whether to pool objects and to return previously created type objects consistent with these rules.

If an error occurs, this method shall return a nil value.

Parameter element_type - The type of all objects that can be stored in a sequence of the new type. If this argument is nil, the operation shall fail with `RETCODE_BAD_PARAMETER`.

Parameter `bound` – The maximum number of elements that may be stored in a map of the new type. If this argument is equal to `LENGTH_UNLIMITED`, the sequence type shall be considered to be unbounded.

7.5.2.2.5 Operations: `create_string_type`, `create_wstring_type`

Create and return a new `DynamicTypeBuilder` object representing a string type. The element type of the result returned by `create_string_type` shall be `Char8`. The element type of the result returned by `create_wstring_type` shall be `Char16`.

All string types having equal element types and equal bounds shall be considered equal. An implementation may therefore elect whether to always return a new object from this method or whether to pool objects and to return previously created type objects consistent with these rules.

If an error occurs, this method shall return a nil value.

Parameter `bound` – The maximum number of elements that may be stored in a string of the new type. If this argument is equal to `LENGTH_UNLIMITED`, the string type shall be considered to be unbounded.

7.5.2.2.6 Operation: `create_type`

Create and return a new `DynamicTypeBuilder` object as described by the given type descriptor. This method is the conventional mechanism for creating structured, enumerated, and alias types, although it can also be used to create types of other kinds. All objects returned by this operation should eventually be deleted by calling `delete_type`.

Parameter `descriptor` – The properties of the new type to create. If this argument is nil or inconsistent (as indicated by its `is_consistent` operation), this operation shall fail and return a nil value.

7.5.2.2.7 Operation: `create_type_copy`

Create and return a new `DynamicTypeBuilder` object with a copy of the state of the given type. All objects returned by this operation should eventually be deleted by calling `delete_type`.

Parameter `type` – The initial state of the new type to create. If this argument is nil, this operation shall fail and return a nil value.

7.5.2.2.8 Operation: `create_type_w_type_object`

Create and return a new `DynamicTypeBuilder` object that describes a type identical to that described by the given `TypeObject` object. Subsequent changes to the new `DynamicTypeBuilder` object shall not be reflected in the input `TypeObject` object. All objects returned by this operation should eventually be deleted by calling `delete_type`.

Parameter `type_object` – The initial state of the new type to create.

7.5.2.2.9 Operation: `delete_instance`

Reclaim any resources associated with any object(s) previously returned from `get_instance`. Any references to these objects held by previous callers of this operation may become invalid at the discretion of the implementation.

This operation shall fail with `RETCODE_ERROR` if it fails for any implementation-specific reason.

7.5.2.2.10 Operation: `delete_type`

Delete the given `DynamicType` object, which was previously created by this factory.

Some “deletions” shall always succeed but shall have no observable effect:

- Deletions of `nil`
- Deletions of objects returned by `get_primitive_type`

Parameter `type` – The type to delete. If this argument is an object that was already deleted, and the implementation is able to detect that fact (which is not required), this operation shall fail with `RETCODE_ALREADY_DELETED`. If an implementation-specific error occurs, this method shall fail with `RETCODE_ERROR`.

7.5.2.2.11 Operation: `get_instance`

Return a `DynamicTypeBuilderFactory` instance that behaves like a singleton, although the caller cannot assume pointer equality for the results of multiple calls. The implementation may return the same object every time or different objects at its discretion. However, if it returns different objects, it shall ensure that they behave equivalently with respect to all programming interfaces specified in this document.

Calling this operation is legal even after `delete_instance` has been called. In such a case, the implementation shall recreate or restore the state of the “singleton” as necessary in order to return a valid object to the caller.

If an error occurs, this method shall return a `nil` value.

7.5.2.2.12 Operation: `get_primitive_type`

Retrieve a `DynamicType` object corresponding to the indicated primitive type kind.

The memory management regime underlying this method is unspecified. Implementations may return references to pre-created objects, they may return new objects with every invocation, or they may take an intermediate approach (for example, lazily creating but then caching objects). Whatever the implementation, the following invariants shall hold:

If an error occurs, this method shall return a `nil` value.

Parameter `kind` – The kind of the primitive type whose representation is to be returned. If the given kind does not correspond to a primitive type, the operation shall fail and return a `nil` value.

7.5.2.2.13 Operation: `create_type_w_uri`

Create and return a new `DynamicType` object by parsing the type description at the given URL.

Applications shall be able to reclaim resources associated with the type returned by this method by calling `delete_type`, just as if the resultant type was created by one of the `create` methods of this class.

If an error occurs, this method shall return a nil value.

Parameter `document_url` – A URL that indicates a type description document, which shall be parsed to create the `DynamicType` object. Implementations shall minimally support the `file://` URL scheme and may support additional schemes. Implementations shall minimally support the XML Type Description format for loaded documents and may support additional Type Descriptions. (Implementations are recommended to provide a tool or other means of translating among their supported Type Representations.)

Parameter `type_name` – The fully qualified name of the type to be loaded from the document that is the target of the URL. If no type exists of this name in the document (which will trivially be the case if the name is nil or the empty string), the operation shall fail and return a nil result.

Parameter `include_paths` – A collection of URLs to directories to be searched for additional type description documents that may be included, directly or indirectly, by the document that is the target of `document_url`. The directory in which the target of `document_url` resides shall be considered on the inclusion search path implicitly and need not be included in this collection. Implementations shall minimally support the `file:` URL scheme and may support additional schemes.

7.5.2.2.14 Operation: `create_type_w_document`

Create and return a new `DynamicType` object by parsing the type description contained in the given string.

Applications shall be able to reclaim resources associated with the type returned by this method by calling `delete_type`, just as if the resultant type was created by one of the `create` methods of this class.

If an error occurs, this method shall return a nil value.

Parameter `document` – A type description document, which shall be parsed to create the `DynamicType` object. Implementations shall minimally support the XML Type Description format for loaded documents and may support additional Type Descriptions. (Implementations are recommended to provide a tool or other means of translating among their supported Type Representations.)

Parameter `type_name` – The fully qualified name of the type to be loaded from the `document`. If no type exists of this name in the document (which will trivially be the case if the name is nil or the empty string), the operation shall fail and return a nil result.

Parameter `include_paths` – A collection of URLs to directories to be searched for additional type description documents that may be included, directly or indirectly, by the `document`

argument. Implementations shall minimally support the `file://` URL scheme and may support additional schemes.

7.5.2.3 AnnotationDescriptor

An `AnnotationDescriptor` packages together the state of an annotation as it is applied to some element (not an annotation type). `AnnotationDescriptor` objects have value semantics, allowing them to be deeply copied and compared.

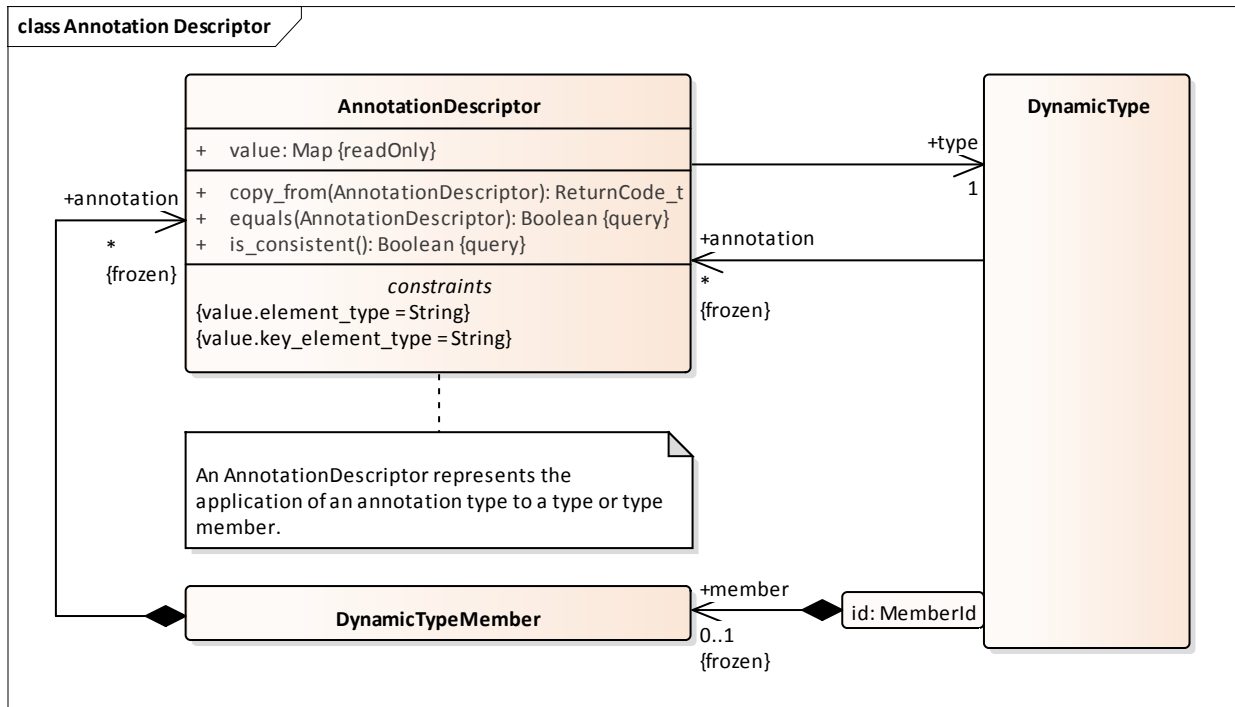


Figure 27 – Annotation Descriptor

Table 49 – AnnotationDescriptor properties and operations

<i>AnnotationDescriptor</i>		
Properties		
type	DynamicType	
value	Map<String<Char8,256>, String<Char8,256>>	
Operations		
copy_from		ReturnCode_t
	other	AnnotationDescriptor
equals		Boolean
	other	AnnotationDescriptor
is_consistent		Boolean

7.5.2.3.1 Operation: `copy_from`

Overwrite the contents of this descriptor with those of another descriptor such that subsequent calls to `equals`, passing the same argument as to this method, return `true`. The other descriptor shall not be changed by this operation.

If this operation fails in an implementation-specific way, this operation shall return `RETCODE_ERROR`.

Parameter `other` – The descriptor whose contents are to be copied. If this argument is `nil`, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.3.2 Operation: `equals`

Two annotation descriptors `ad1` and `ad2` are considered equal if and only if all of the following apply:

- Their type properties refer to equal types.
- For every string `s1` for which `ad1.value[s1]` does not exist, `ad2.value[s1]` also does not exist.
- For every string `s1` for which `ad2.value[s1]` does not exist, `ad1.value[s1]` also does not exist.
- For every string `s1` for which `ad1.value[s1]` is a non-nil string `ad1-s2`, `ad2.value[s1]` is a non-nil string `ad2-s2` such that `ad1-s2` equals `ad2-s2`.
- For every string `s1` for which `ad2.value[s1]` is a non-nil string `ad2-s2`, `ad1.value[s1]` is a non-nil string `ad1-s2` such that `ad1-s2` equals `ad2-s2`.

Parameter `other` – Another descriptor to compare to this descriptor. If this argument is `nil`, this operation shall return `false`.

7.5.2.3.3 Operation: `is_consistent`

Indicate whether this descriptor describes a valid annotation type instantiation. An annotation descriptor is considered consistent if and only if all of the following qualities apply:

- The `type` property refers to a non-nil type of kind `ANNOTATION_TYPE`.
- For every pair of strings `s1` and `s2` such that `value[s1]` equals `value[s2]`:
 - String `s1` is the name of an attribute defined by the annotation type referred to by the `type` property.
 - String `s2` is a well-formed string representation of an object of the type of the attribute named by `s1`.

7.5.2.3.4 Property: `type`

The `type` property contains a reference to the annotation type, of which this descriptor describes an instantiation.

When an annotation descriptor is newly created, this reference shall be nil.

7.5.2.3.5 Property: value

This property contains a mapping from the names of attributes defined by `type` to valid values of that type. Any attribute defined by `type` but for which no name appears in this property shall be considered to have its default value.

Every attribute value in this property is represented as a string although annotation type members can have other types as well. A string representation of a data value is considered well-formed if it would be a valid IDL literal of the corresponding type with the following qualifications:

- String and character literals shall not be surrounded by quotation characters (“” or ‘’).
- All expressions shall be fully evaluated such that no operators or other non-literal characters occur in the value. For example, “5” shall be considered a well-formed string representation of the integer quantity *five*, but “2 + ENUM_VALUE_THREE” shall not be.

7.5.2.4 TypeDescriptor

A `TypeDescriptor` packages together the state of a type. `TypeDescriptor` objects have value semantics, allowing them to be deeply copied and compared.

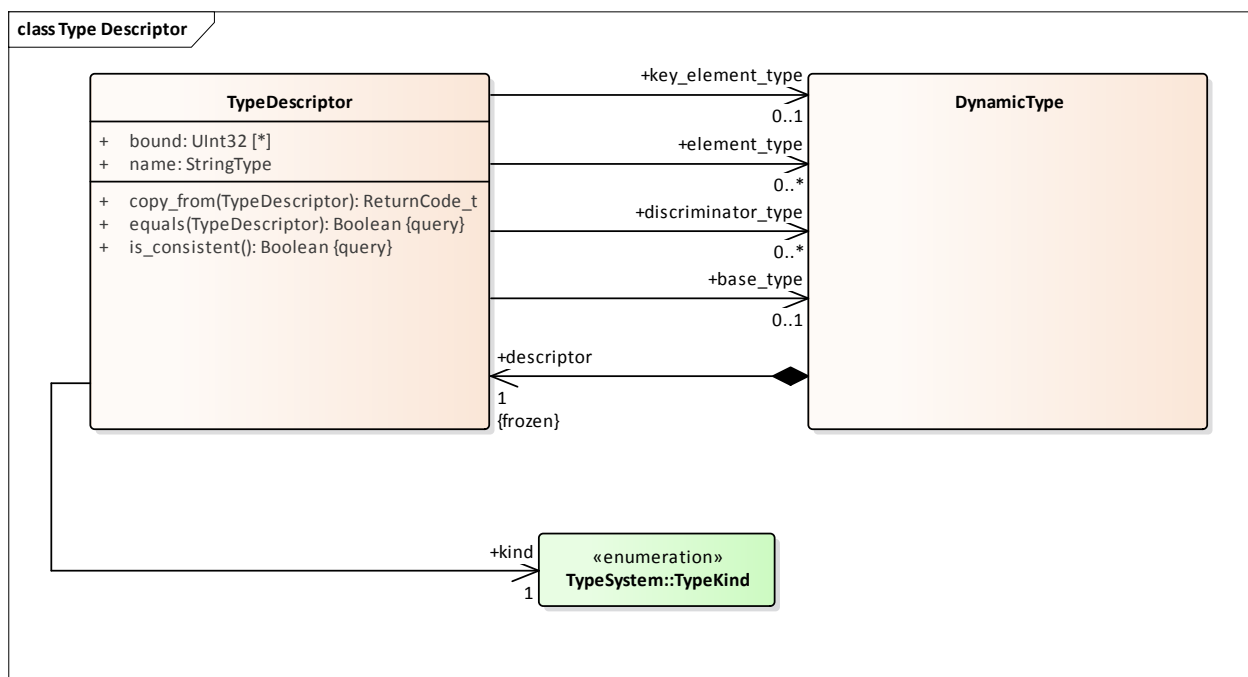


Figure 28 – Type Descriptor

Table 50 – TypeDescriptor properties and operations

<i>TypeDescriptor</i>		
Properties		
kind	TypeKind	
name	string<Char8,256>	
base_type	DynamicType [0..1]	
discriminator_type	DynamicType [0..1]	
bound	UInt32 [*]	
element_type	DynamicType [0..1]	
key element type	DynamicType [0..1]	
Operations		
copy_from		ReturnCode t
	other	TypeDescriptor
equals		Boolean
	other	TypeDescriptor
is consistent		Boolean

7.5.2.4.1 Property: base_type

Another type definition, on which the type described by this descriptor is based. Specifically:

- If this descriptor represents a structure type, `base_type` indicates the supertype of that type. A nil value of this property indicates that the structure type has no supertype.
- If this descriptor represents an alias type, `base_type` indicates the type being aliased. A nil value for this property is not considered consistent.

In all other cases, a consistent descriptor shall have a nil value for this property.

7.5.2.4.2 Property: bound

The `bound` property indicates the bound of collection and similar types.

- If this descriptor represents an array type, the length of the property value indicates the number of dimensions in the array, and each value indicates the bound of the corresponding dimension.
- If this descriptor represents a sequence, map, bitmask, or string type, the length of the property value is one and the integral value in that property indicates the bound of the collection.

In all other cases, a consistent descriptor shall have a nil value for this property.

7.5.2.4.3 Operation: `copy_from`

Overwrite the contents of this descriptor with those of another descriptor such that subsequent calls to `equals`, passing the same argument as to this method, return `true`. The other descriptor shall not be changed by this operation.

If this operation fails in an implementation-specific way, this operation shall return `RETCODE_ERROR`.

Parameter `other` – The descriptor whose contents are to be copied. If this argument is `nil`, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.4.4 Property: `discriminator_type`

If this descriptor represents a union type, `discriminator_type` indicates the type of the discriminator of the union. It must not be `nil` for the descriptor to be consistent.

If this descriptor represents any other kind of type, this property must be `nil` for this descriptor to be consistent.

7.5.2.4.5 Property: `element_type`

If this descriptor represents an array, sequence, or string type, this property indicates the element type of the collection. It must not be `nil` for the descriptor to be consistent.

If this descriptor represents a map type, this property indicates the *value* element type of the map. It must not be `nil` for the descriptor to be consistent.

If this descriptor represents a bitmask type, this property must indicate a Boolean type for the descriptor to be consistent.

If this descriptor represents any other kind of type, this property must be `nil` for the descriptor to be consistent.

7.5.2.4.6 Operation: `equals`

Two type descriptors are considered equal if and only if the values of all of the properties identified in Table 50 above are equal in each of them.

Parameter `other` – Another descriptor to compare to this one. If this argument is `nil`, the operation shall return `false`.

7.5.2.4.7 Operation: `is_consistent`

Indicates whether the states of all of this descriptor's properties are consistent. The definitions of consistency for each property are given in the clause corresponding to that property.

7.5.2.4.8 Property: `key_element_type`

If this descriptor represents a map type, this property indicates the *value* element type of the map. It must not be `nil` for the descriptor to be consistent.

If this descriptor represents any other kind of type, this property must be nil for the descriptor to be consistent.

7.5.2.4.9 Property: `kind`

An enumerated value that indicates what “kind” of type this descriptor describes: a structure, a sequence, etc.

7.5.2.4.10 Property: `name`

The fully qualified name of the type described by this descriptor. To be consistent, this name must be a valid identifier for the given type kind, as defined elsewhere in this document.

7.5.2.5 MemberId

The type `MemberId` is an alias to `UInt32` and is used for the purpose of representing the ID of a member of a structured type.

It is also used to type the constant `MEMBER_ID_INVALID`, which is a sentinel indicating a member ID that is missing, irrelevant, or otherwise invalid in a given context.

7.5.2.6 DynamicTypeMember

A `DynamicTypeMember` represents a “member” of a type. A “member” in this sense may be a member of an aggregated type, a constant within an enumeration, or some other type substructure. Specifically, the behavior is as described in Table 51 below based on the `TypeKind` of the `DynamicType` to which the member belongs.

Table 51 – DynamicMember behavior

<i>Type Kind</i>	<i>Meaning</i>
ANNOTATION_TYPE	For these aggregated types, a “member” in this sense has the same meaning as it does in the definition of aggregated types generally.
STRUCTURE_TYPE	
UNION_TYPE	
BITMASK_TYPE	Each named flag in a bitmask shall be considered to be a “member” of that bitmask with <code>Boolean</code> type.
ENUMERATION_TYPE	Each literal in the enumeration shall be considered a “member” of the type. These members shall have the type of the enclosing enumeration itself.
ALIAS_TYPE	The behavior is as it would be for the alias’s base type.

No other type kinds are considered to have members.

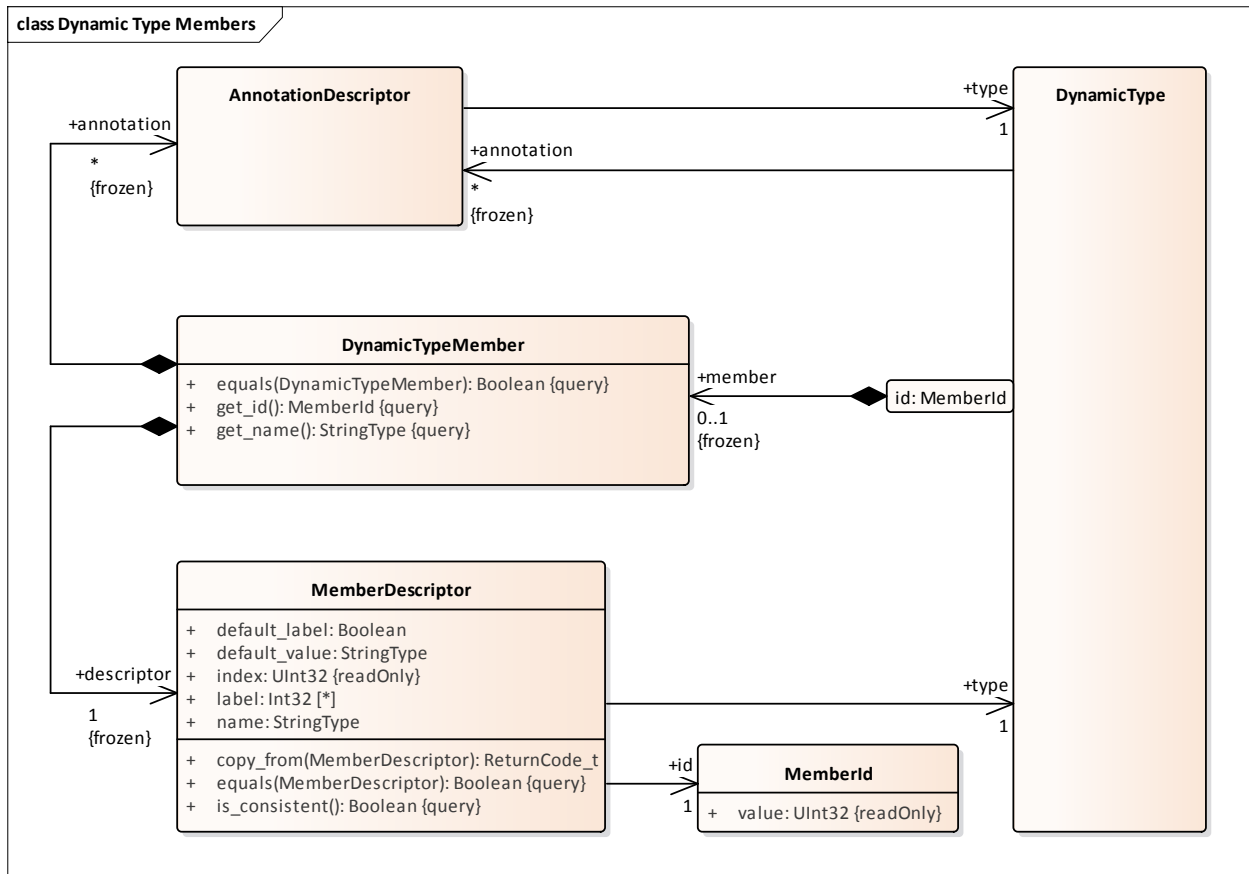


Figure 29 – Dynamic Type Members

`DynamicTypeMember` objects have reference semantics; however, there is an `equals` operation to allow them to be deeply compared.

Table 52 – `DynamicTypeMember` properties and operations

<i>DynamicTypeMember</i>		
Properties		
annotation		read-only AnnotationDescriptor [*]
Operations		
get_descriptor		DDS::ReturnCode_t
	inout descriptor	MemberDescriptor
equals		Boolean
	other	DynamicTypeMember
get_name		string<Char8,256>
get_id		MemberId

7.5.2.6.1 Property: `annotation`

This property provides all annotations previously applied to this member.

7.5.2.6.2 Operation: `get_descriptor`

This operation provides a summary of the state of this type. It overwrites the state of an application-provided object.

If the argument is `nil`, this operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.6.3 Operation: `equals`

Two members shall be considered equal if and only if they belong to the same type and all of their respective properties, as identified in Table 52 above, are equal.

7.5.2.6.4 Operation: `get_id`

This convenience operation provides the member ID of this member. Its result shall be identical to the ID value that is a member of the `descriptor` property.

7.5.2.6.5 Operation: `get_name`

This convenience operation provides the name of this member. Its result shall be identical to the name string that is a member of the `descriptor` property.

7.5.2.7 MemberDescriptor

A `MemberDescriptor` packages together the state of a `DynamicTypeMember`. `MemberDescriptor` objects have value semantics, allowing them to be deeply copied and compared.

Table 53 – MemberDescriptor properties and operations

<i>MemberDescriptor</i>	
Properties	
<code>name</code>	<code>String<Char8,256></code>
<code>id</code>	<code>MemberId</code>
<code>type</code>	<code>DynamicType</code>
<code>default value</code>	<code>string</code>
<code>index</code>	<code>read-only UInt32</code>
<code>label</code>	<code>Int64 [*]</code>
<code>default label</code>	<code>Boolean</code>

Operations		
copy_from		ReturnCode_t
	other	MemberDescriptor
equals		Boolean
	other	MemberDescriptor
is_consistent		Boolean

7.5.2.7.1 Operation: copy_from

Overwrite the contents of this descriptor with those of another descriptor such that subsequent calls to `equals`, passing the same argument as to this method, return `true`. The other descriptor shall not be changed by this operation.

If this operation fails in an implementation-specific way, this operation shall return `RETCODE_ERROR`.

Parameter `other` – The descriptor whose contents are to be copied. If this argument is `nil`, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.7.2 Property: default_label

For this descriptor to be consistent, this property must be `true` if this descriptor identifies the default member of a union type or `false` if not. A default union member may have additional explicit labels (indicated in the `label` property), but these are semantically irrelevant, as the default member would be in effect or not regardless of their presence or absence.

7.5.2.7.3 Property: default_value

This property provides the member’s default value in string form. A string representation of a data value is considered well-formed if it would be a valid IDL literal of the corresponding type with the following qualifications:

- String and character literals shall not be surrounded by quotation characters (“” or “”).
- All expressions shall be fully evaluated such that no operators or other non-literal characters occur in the value. For example, “5” shall be considered a well-formed string representation of the integer quantity *five*, but “2 + ENUM_VALUE_THREE” shall not be.

A `nil` or empty string indicates that the member takes the “default default” value for its type. This rule shall *always* be used when the member is of a type for which IDL provides no syntax to express a literal value (for example, structures or maps) and *may* be used for any other type.

Design rationale: An instance of `DynamicData` might have been used here as an alternative. However, since every default literal can be expressed as a string anyway (i.e., as it is in IDL), and string objects are expected to be more lightweight than `DynamicData` implementations, that representation was preferred.

7.5.2.7.4 Operation: equals

Two descriptors are considered equal if and only if the values of all of the properties identified in Table 53 above are equal in each of them.

Parameter *other* – Another descriptor to compare to this one. If this argument is nil, the operation shall return *false*.

7.5.2.7.5 Property: id

If this member belongs to an aggregated type, this property indicates the member's ID.

- When a descriptor is used to add a new member to a type, this property may be set to `MEMBER_ID_INVALID`; in that case, the implementation shall select an ID for the new member that is one more than the current maximum member ID in the type. If the value of this property is *not* `MEMBER_ID_INVALID`, it must be set to a value within a legal range.
- When a descriptor is retrieved from an existing member, this property shall reflect the actual ID of the member. It shall therefore not be `MEMBER_ID_INVALID`, and it shall fall within a legal range.

If this member does not belong to an aggregated type, this property *must* be `MEMBER_ID_INVALID`, or the descriptor is not consistent.

7.5.2.7.6 Property: index

This property indicates the order of definition of this member within its type, relative to the type's other members. The first member shall have index zero, the next one, and so on.

When a descriptor is used to add a new member to a type, any value greater than the current largest index value in the type shall be taken to indicate that the new member will become the last member, whatever the index; member indices within a type shall not be discontinuous. Alternatively, if this property is set to an index at which a member already exists, that member and all those after it shall be shifted up by a single index value to make room for the new member.

When a descriptor is retrieved from an existing member, this property shall always reflect the actual index at which the member exists.

7.5.2.7.7 Operation: is_consistent

A descriptor shall be considered consistent if and only if all of the values of its properties are considered consistent. The meaning of consistency for each of these is defined here in the appropriate clause.

7.5.2.7.8 Property: label

If the type to which the member belongs is a union, this property indicates the case labels that apply to this member. If `default_label` is false, it must not be empty. In addition, no two members of the same union can specify the same label value.

If the type to which the member belongs is *not* a union, this property's value must be empty to be consistent.

7.5.2.7.9 Property: name

This property indicates the name of this member. The value must be a well-formed member name.

7.5.2.7.10 Property: type

This property indicates the type of the member's value. It must not be nil and must indicate a type that can legally type a member according to the Type System Model.

7.5.2.8 DynamicType

A `DynamicType` object represents a particular type defined according to the Type System. `DynamicType` objects have reference semantics because of the large number of references to them that are expected to exist (e.g., in each `DynamicData` object created from a given `DynamicType`). However, the type nevertheless provides operations to allow copying and comparison by value.




Figure 30 – Dynamic Type

Table 54 – DynamicType properties and operations

<i>DynamicType</i>	
Properties	
member_by_name	read-only string<Char8,256> → DynamicTypeMember [0..1]
member	read-only MemberId → DynamicTypeMember [0..1]
annotation	read-only AnnotationDescriptor [*]

Operations		
get_descriptor		DDS::ReturnCode_t
	inout_descriptor	TypeDescriptor
equals		Boolean
	other	DynamicType
get_name		string<Char8,256>
get_kind		TypeKind

7.5.2.8.1 Property: annotation

This property provides all annotations that have previously been applied to this type.

7.5.2.8.2 Operation: get_descriptor

This operation provides a summary of the state of this type. It overwrites the state of an application-provided object.

If the argument is nil, this operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.8.3 Operation: equals

Two types shall be considered equal if and only if all of their respective properties, as identified in Table 54 above, are equal.

7.5.2.8.4 Operation: get_kind

This convenience operation indicates the kind of this type (e.g., integer, structure, etc.). Its result shall be the same as the kind indicated by the type's `descriptor` property.

7.5.2.8.5 Operation: get_name

This convenience operation provides the fully qualified name of this type. It shall be identical to the name string that is a member of the `descriptor` property.

7.5.2.8.6 Property: member_by_name

This property contains a mapping from the name of a member of this type to the member itself. As described in Table 55 below, not only members of aggregated types are considered “members” here: the constituents of enumerations, bitmasks, and other kinds of types are also considered to be “members” for the purposes of this property.

Table 55 – DynamicType::member_by_name behavior

<i>Type Kind</i>	<i>Behavior</i>
ANNOTATION_TYPE	The member descriptor must describe a consistent annotation type member. If the descriptor does not satisfy these constraints, the operation shall fail with <code>RETCODE_BAD_PARAMETER</code> .

ALIAS_TYPE	The behavior is as it would be for the alias's base type. If adding a member is not defined for the alias's base type, this operation shall fail with <code>RETCODE_PRECONDITION_NOT_MET</code> .
BITMASK_TYPE	The member descriptor must describe a Boolean flag with a value within the bound of this bitmask type. If the descriptor does not satisfy these constraints, the operation shall fail with <code>RETCODE_BAD_PARAMETER</code> .
ENUMERATION_TYPE	The member descriptor must describe a literal with the type of this enumeration. If the descriptor does not satisfy these constraints, the operation shall fail with <code>RETCODE_BAD_PARAMETER</code> .
STRUCTURE_TYPE	The member descriptor must describe a consistent structure member. If the descriptor does not satisfy this constraint, the operation shall fail with <code>RETCODE_BAD_PARAMETER</code> .
UNION_TYPE	The member descriptor must describe a consistent union member. If the descriptor does not satisfy this constraint, the operation shall fail with <code>RETCODE_BAD_PARAMETER</code> .

The lifecycle of a `DynamicTypeMember` object is governed by that of the `DynamicType` that contains it. The former shall be considered to exist logically from the time the corresponding member is added to the latter and until such time as the latter is deleted. Implementations may allocate and de-allocate `DynamicTypeMember` objects more frequently, provided that:

- Users of the `DynamicTypeMember` class are not required to explicitly delete objects of that class.
- Changes to one `DynamicTypeMember` object representing a given member shall be reflected in all observable `DynamicTypeMember` objects representing the same member.
- All `DynamicTypeMember` objects representing the same member shall compare as equal according to their `equals` operations.

7.5.2.8.7 Property: `member`

This property contains a mapping from the member ID of a member of this (aggregated) type to the member itself.

- If this type is an aggregated type, the collection of members available through this property shall be equal to (element order notwithstanding) that available through the `member_by_name` property.
- If this type is *not* an aggregated type, the collection of members available through this property shall be empty.

7.5.2.9 `DynamicTypeBuilder`

A `DynamicTypeBuilder` object represents a transitional state of a particular type defined according to the Type System. It is used to instantiate concrete `DynamicType` objects.

Table 56 – DynamicTypeBuilder properties and operations

<i>DynamicTypeBuilder</i>		
Properties		
member_by_name	read-only string<Char8,256> → DynamicTypeMember [0..1]	
member	read-only MemberId → DynamicTypeMember [0..1]	
annotation	read-only AnnotationDescriptor [*]	
Operations		
get_descriptor		DDS::ReturnCode_t
	inout descriptor	TypeDescriptor
equals		Boolean
	other	DynamicType
get_name		string<Char8,256>
get_kind		TypeKind
add_member		ReturnCode_t
	descriptor	MemberDescriptor
apply_annotation		ReturnCode_t
	descriptor	AnnotationDescriptor
apply_annotation_to_member		ReturnCode_t
	member_id	MemberId
	descriptor	AnnotationDescriptor
build		DynamicType

7.5.2.9.1 Operation: add_member

Add a “member” to this type, where the new “member” has the meaning defined in the specification of the `DynamicTypeMember` class. Specifically, the behavior shall be as described in Table 55 in Clause 7.5.2.8.6, “Property: member_by_name”. For type kinds not given in that table, this operation shall fail with `RETCODE_PRECONDITION_NOT_MET`.

Following a successful return, the new member shall appear in the member property and possibly in the `member_by_id` property, based on the definition of that property.

Parameter `descriptor` – A descriptor of the new member to be added. If this argument is nil, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.9.2 Property: `annotation`

This property provides all annotations that have previously been applied to this type with `apply_annotation`.

7.5.2.9.3 Operation: `apply_annotation`

Apply the given annotation to this type. It shall subsequently appear in the `annotation` property.

Parameter `descriptor` – A consistent descriptor for the annotation to apply. If this argument is not consistent, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.9.4 Operation: `apply_annotation_to_member`

Apply the given annotation to this member. It shall subsequently appear in the `annotation` property of the identified member.

Parameter `member_id` – Identifies the member to which the annotation shall be applied.

Parameter `descriptor` – A consistent descriptor for the annotation to apply. If this argument is not consistent, the operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.9.5 Operation: `build`

Create an immutable `DynamicType` object containing a snapshot of this builder's current state. Subsequent changes to this builder, if any, shall have no observable effect on the states of any previously created `DynamicTypes`.

7.5.2.9.6 Operation: `get_descriptor`

This operation provides a summary of the state of this type. It overwrites the state of an application-provided object.

If the argument is `nil`, this operation shall fail with `RETCODE_BAD_PARAMETER`.

7.5.2.9.7 Operation: `equals`

Two types shall be considered equal if and only if all of their respective properties, as identified in Table 56 above, are equal.

7.5.2.9.8 Operation: `get_kind`

This convenience operation indicates the kind of this type (e.g., integer, structure, etc.). Its result shall be the same as the kind indicated by the type's descriptor property.

7.5.2.9.9 Operation: `get_name`

This convenience operation provides the fully qualified name of this type. It shall be identical to the name string that is a member of the descriptor property.

7.5.2.9.10 Property: `member_by_name`

This property contains a mapping from the name of a member of this type to the member itself. As described in the case of `add_member`, not only members of aggregated types are considered

“members” here: the constituents of enumerations, bitmasks, and other kinds of types are also considered to be “members” for the purposes of this property.

The lifecycle of a `DynamicTypeMember` object is governed by that of the `DynamicTypeBuilder` that contains it. The former shall be considered to exist logically from the time the corresponding member is added to the latter and until such time as the latter is deleted. Implementations may allocate and de-allocate `DynamicTypeMember` objects more frequently, provided that:

- Users of the `DynamicTypeMember` class are not required to explicitly delete objects of that class.
- Changes to one `DynamicTypeMember` object representing a given member shall be reflected in all observable `DynamicTypeMember` objects representing the same member.
- All `DynamicTypeMember` objects representing the same member shall compare as equal according to their equals operations.

7.5.2.9.11 Property: `member`

This property contains a mapping from the member ID of a member of this (aggregated) type to the member itself.

- If this type is an aggregated type, the collection of members available through this property shall be equal to (element order notwithstanding) that available through the `member_by_name` property.
- If this type is not an aggregated type, the collection of members available through this property shall be empty.

7.5.2.10 `DynamicDataFactory`

This class is logically a singleton (although it need not technically be a singleton in practice). Its “only” instance is the starting point for creating and deleting `DynamicData` and objects, just like the singleton `DomainParticipantFactory` is the starting point for creating `DomainParticipant` objects.

Table 57 – `DynamicDataFactory` properties and operations

<i>DynamicDataFactory</i>		
Operations		
<code>static get instance</code>		<code>DynamicDataFactory</code>
<code>static delete instance</code>		<code>ReturnCode t</code>
<code>create_data</code>		<code>DynamicData</code>
	<code>type</code>	<code>DynamicType</code>
<code>delete_data</code>		<code>ReturnCode t</code>
	<code>data</code>	<code>DynamicData</code>

7.5.2.10.1 Operation: `create_data`

Create and return a new data sample. All objects returned by this operation should eventually be deleted by calling `delete_data`.

Parameter `type` - The type of the sample to create.

7.5.2.10.2 Operation: `delete_data`

Dispose of a data sample, reclaiming any associated resources.

Parameter `data` - The data sample to delete.

7.5.2.10.3 Operation: `delete_instance`

Reclaim any resources associated with the object(s) previously returned from `get_instance`. Any references to these objects held by previous callers may become invalid at the implementation's discretion.

This operation shall return `RETCODE_ERROR` if it fails for any implementation-specific reason.

7.5.2.10.4 Operation: `get_instance`

Return a `DynamicDataFactory` instance that behaves like a singleton, although callers cannot assume pointer equality across invocations of this operation. The implementation may return the same object every time or different objects at its discretion. However, if it returns different objects, it shall ensure that they behave equivalently with respect to all programming interfaces specified in this document.

It is legal to call this operation even after `delete_instance` has been called. In such a case, the implementation shall recreate or restore the “singleton” as necessary to ensure that it can return a valid object to the caller.

If an error occurs, this method shall return a nil value.

7.5.2.11 `DynamicData`

Each object of the `DynamicData` class represents a corresponding object of the type represented by the `DynamicData` object's `DynamicType`.

`DynamicData` objects have reference semantics; however, there is an `equals` operation to allow them to be deeply compared.

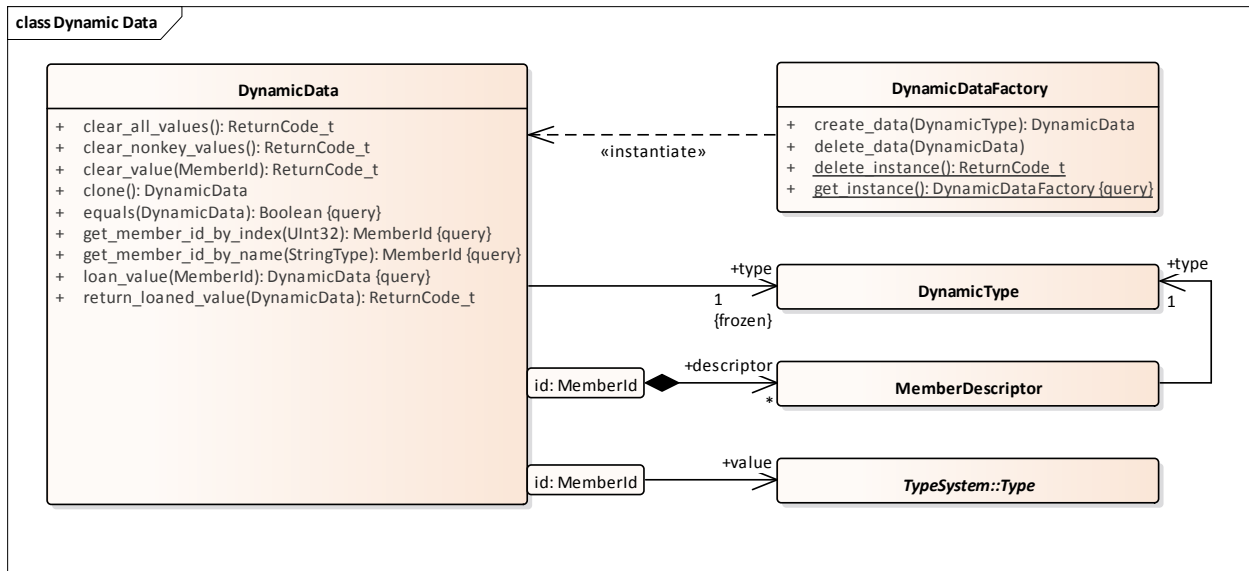


Figure 31 – Dynamic Data and Dynamic Data Factory

Table 58 below summarizes the properties and operations supported by `DynamicData` objects.

Table 58 – `DynamicData` properties and operations

<i>DynamicData</i>		
Properties		
value	MemberId	→ Type [0..1]
type	read-only	DynamicType
descriptor	MemberId	→ MemberDescriptor
Operations		
get_member_id_by_name		MemberId
	name	string<Char8,256>
get_member_id_at_index		MemberId
	index	UInt32
get item count		UInt32
equals		Boolean
	other	DynamicData
clear all values		ReturnCode_t
clear nonkey values		ReturnCode_t
clear_value		ReturnCode_t
	id	MemberId

loan_value		DynamicData
	member_id	MemberId
return_loaned_value		ReturnCode_t
	value	DynamicData
clone		DynamicData

7.5.2.11.1 Property: value; Operations: get_member_id_by_name and get_member_id_at_index

Many of the properties and operations defined by this class refer to values within the sample, which are identified by name, member ID, or index. What constitutes a value within a sample, and which means of accessing it are valid, depends on the type of this sample.

- If this object is of an aggregated type, values correspond to the type’s members and can be accessed by name, member ID, or index.
- If this object is of a sequence or string type, values correspond to the elements of the collection. These elements must be accessed by index; the mapping from index to member ID is unspecified.
- If this object is of a map type, values correspond to the values of the map. Map keys are implicitly converted to strings and can thus be used to look up map values by name. Map values can also be accessed by index, although the order is unspecified.
- If the object is of an array type, values correspond to the elements of the array. These elements must be accessed by index; the mapping from index to member ID is unspecified. If the array is multi-dimensional, elements are accessed as if they were “flattened” into a single-dimensional array in the order specified by the IDL specification.
- If the object is of a bitmask type, values correspond to the flags within the bitmask and are all of `Boolean` type. Named flags can be accessed using that name; any bit within the bound of the bitmask may be accessed by its index. The mappings from name and index to member ID are unspecified.
- If the object is of an enumeration or primitive type, it has no contained values. However, the value of the sample itself may be indicated by “name” using a nil or empty string, by “ID” by passing `MEMBER_ID_INVALID`, or by “index” by passing index 0.

Note that indices used here are always relative to other values *in a particular DynamicData object*. Even though member *definitions* within aggregated types have a well-defined order, the same is *not* true within data samples or across data samples. Specifically, the index at which a member of an aggregated type appears in a particular data sample may not match that in which it appears in the corresponding type and may not match the index at which it appears in a different data sample. There are several reasons for these inconsistencies:

- The producer of the sample may be using a slightly different variant of the type than the consumer, which may add to, or omit elements from, the set of members known to the consumer.
- An optional member may have no value; in such a case, it will be omitted, thereby decreasing the index of every subsequent member.
- A non-optional member may likewise be omitted (which semantically is equivalent to it taking its default value). An implementation may discretionarily omit such members (e.g., to save space).
- Preserving member order is not necessary or even desirable (e.g., for performance reasons) for certain data representations.

7.5.2.11.2 Property: `descriptor`

This property shall contain a descriptor for each value in this object, identified by the member ID. The meaning of the member ID shall be as it is described for the `value` property.

7.5.2.11.3 Clearing Values: Operations `clear_value`, `clear_all_values`, and `clear_nonkey_values`

The meaning of “clearing” a member depends on the type of data represented by this sample:

- If this sample is of an aggregated type, and the indicated member is optional, remove it. If the indicated member is not optional, set it to its default value.
- If this sample is of a variable-length collection type, remove the indicated element, shifting any subsequent elements to the next-lowest index.
- If the sample is of an array type, set the indicated element to its default value.
- If the sample is of a bitmask type, clear the indicated bit.
- If the sample is of an enumerated type, set it to the first value of the enumerated type.
- If the sample is of a primitive type, set it to its default value.

The `clear_all_members` takes the above action for each value in turn. The `clear_nonkey_value` operation has exactly the same effect as `clear_all_values` with one exception: the values of key fields of aggregated types retain their values.

7.5.2.11.4 Operation: `clone`

Create and return a new data sample with the same contents as this one. A comparison of this object and the clone using `equals` immediately following this call will return `true`.

7.5.2.11.5 Operation: `equals`

Two data samples are considered to be equal if and only if all of the following conditions hold:

- Their respective type definitions are equal.

- All contained values are equal and occur in the same order.
- If the samples' type is an aggregated type, the previous rule shall be amended as follows:
 - Members shall be compared without regard to their order.
 - One of the samples may omit a non-optional member that is present in the other if that member takes its default value in the latter sample.

7.5.2.11.6 Operation: `get_item_count`

The “item count” of the data depends on the type of the object.

- *If the object is of a collection type*, return the number of elements currently in the collection. In the case of an array type, this value will always be equal to the product of the bounds of all array dimensions.
- *If the object is of a bitmask type*, return the number of named flags that are currently set in the bitmask.
- *If the object is of a structure or annotation type*, return the number of members in the object. This value may be different than the number of members in the corresponding `DynamicType`—for example, some optional members may be omitted.
- *If the object is of a union type*, return the number of members in the object. This value will always be two: the discriminator and the current member corresponding to it.
- *If the object is of a primitive or enumerated type*, it is atomic: return one.
- *If the object is of an alias type*, return the value appropriate for the alias's base type.

7.5.2.11.7 Operations: `loan_value` and `return_loaned_value`

The “loan” operations loan to the application a `DynamicData` object representing a value within this sample. These operations allow applications to visit values without allocating additional `DynamicData` objects or copying values. This loan shall be returned by the `return_loaned_value` operation.

A given `DynamicData` object may support only a single outstanding loan at a time. That is, after calling a “loan” operation, an application must subsequently call `return_loaned_value` before calling a loan operation again. If an application violates this constraint, the loan operation shall return a nil value.

A loan operation shall also return a nil value if the indicated value does not exist.

The `return_loaned_value` operation shall return `RETCODE_PRECONDITION_NOT_MET` if the provided sample object does not represent an outstanding loan from the sample on which the operation is invoked.

7.5.2.11.8 Property: `type`

This property provides the type that defines the values within this sample. Its value shall not be nil.

7.5.2.11.9 Platform-Specific Model: IDL

The programming language-specific APIs for the Dynamic Type and Dynamic Data classes and their companion classes shall be based on the following IDL definitions, transformed according to the IDL language mapping described above, as expanded below.

The conceptual model refers to the type `Object`, objects of which may be of any concrete type supported by the Type System defined by this specification. The mapping to IDL below represents this multiplicity of concrete types by multiplying the methods implied by the properties, qualifying each method with a concrete type. For example, a qualified association `foo: Int32 → Object` would expand to `get_int32_foo`, `get_int16_foo`, etc. Specifically, the mapping uses the following type expansions:

- Each primitive type has its own expansion. Primitive types can be implicitly promoted to larger primitive types as defined below.
- Strings of `Char8` and `Char16` elements have their own expansions qualified by “string” and “wstring” respectively.
- Enumerated types shall be implicitly converted to any signed integer type having at least as many bits as the enumerated type’s `@bit_bound`. They are thus accessible through those primitive methods.
- Bitmasks shall be implicitly converted to any unsigned integer type having at least as many bits as the bitmask’s `@bit_bound`. They are thus accessible through those primitive methods.
- Alias types shall be implicitly converted to their ultimate base type and are thus accessible through the methods appropriate for that type.
- Sequences of primitive types and strings have their own expansions in which the name of the property has been made plural. Arrays shall also be accessible through these methods.
- Expansions that operate on `DynamicData` objects, qualified by “complex,” catch the remaining cases and offer an alternative approach to accessing values of any of the above types.

If a `DynamicData` object represents an object of a resizable collection type (string, sequence, or map), these setters may also be used to append new elements to the collection.

- For a string or sequence type, use `get_member_id_at_index` to obtain an ID for the index one greater than the current length.
- For a map type, use `get_member_id_by_name` to obtain an ID for the new map key.

As mentioned above, it shall be possible to implicitly promote integral types. These shall be supported during both “get” and “set” operations such that a smaller type promotes to a large type but not vice versa. For example, it shall be possible to get the value of a short integer field as if it were a long integer, and it shall be possible to set the value of a long integer as if it were a short integer. Specifically, the following promotions shall be supported:

- Int16 → Int32, Int64, Float32, Float64, Float128
- Int32 → Int64, Float64, Float128
- Int64 → Float128
- UInt16 → Int32, Int64, UInt32, UInt64, Float32, Float64, Float128
- UInt32 → Int64, UInt64, Float64, Float128
- UInt64 → Float128
- Float32 → Float64, Float128
- Float64 → Float128
- Float128 → (*none*)
- Char8 → Char16, Int16, Int32, Int64, Float32, Float64, Float128
- Char16 → Int32, Int64, Float32, Float64, Float128
- Byte → (*any*)
- Boolean → Int16, Int32, Int64, UInt16, UInt32, UInt64, Float32, Float64, Float128

The complete IDL representation may be found in “Annex C: Dynamic Language Binding.”

7.6 Use of the Type System by DDS

This clause describes how DDS uses the type system.

7.6.1 Topic Model

A DDS topic exists in two senses of the word:

1. **On the network**, with respect to interoperability: This is the sense in which we say that a reader and a writer share the “same” topic, even though they obtain the topic’s definition independently within their implementations.
2. **In application code**, with respect to portability: Each component that uses a topic creates or looks up a local proxy for that topic.

On the network, a given topic is associated with one or more types. A given writer or reader endpoint belongs to one topic and is associated with one of the types of that topic. If a writer and a reader share the same topic, it is assumed that they are intended to communicate with one another. At that point, the Service evaluates the two endpoints to make sure that they specify consistent types (see Clause 7.6.2.4.2, “Rules for Type Consistency Enforcement”) and compatible QoS (see [DDS]).

Typically, in application code, a topic is associated with a single type (as has always been the case in the [DDS] API)⁶. Therefore, multiple API topics may correspond to (different views of) the same network topic. A given reader or writer endpoint is associated with one of them. See Clause 7.6.3, “Local API Extensions”, for definitions of the programming interfaces that support this polymorphism.

Generic services (e.g., logger, monitor) may discover a topic associated with one or more types. Such services may be able to handle all representations of the types, without ever having type specific knowledge hardcoded into them.

7.6.2 Discovery and Endpoint Matching

The enhanced Type System and the richer set of available Data Representations necessitate extensions to the discovery and endpoint matching process defined by the DDS specification, which may be divided into three categories:

- **Data Representation:** The multiplicity of data representations introduced by this specification creates the possibility that different `DataWriter` and `DataReader` endpoints in a single system may support different combinations of representations. It is therefore necessary to define a mechanism whereby endpoints can inform each other of the representations they support and thereby negotiate communication.
- **Discovery-Time Data Typing:** The dynamic features of this specification depend on the ability of components to discover the data types used by their peers.
- **Type Consistency Enforcement:** One of the criteria for `DataWriter-DataReader` matching defined by DDS is that the type names of each must match exactly. In complex dynamic systems, this restriction can prove overly limiting. Based on the type compatibility rules defined by this specification, matching endpoints shall be permitted to declare types that are not identical but nevertheless have well-defined relationships with one another.

These extensions are defined in the following sections.

7.6.2.1 Data Representation QoS Policy

With multiple standard data Representations available, and vendor-specific extensions possible, `DataWriters` and `DataReaders` must be able to negotiate which data representation(s) to use. This negotiation shall occur based on a new QoS policy: `DataRepresentationQoSPolicy`.

7.6.2.1.1 `DataRepresentationQoSPolicy`: Conceptual Model

The conceptual model for data representation negotiation consists of several parts:

- The identification of data representations.

⁶ **Design rationale (non-normative):** This constraint keeps the programming model the same for both XTypes-supporting and non-XTypes-supporting implementations, and it keeps the mental model simple for the majority of programmers, who will not be aware of the presence of multiple types in their topics.

- The specification of supported and preferred representations by `DataReaders` and `DataWriters`.
- The algorithm by which a suitable representation is chosen for a particular `DataReader/DataWriter` pair, given the supported representations of each.

Each data representation shall be identified by a two-byte signed integer value, the “representation identifier.” Within the range of such a value, the negative values shall be reserved for definition by DDS implementations. The remainder of the range shall be reserved for the OMG for use in future specifications, including this specification.

Within the OMG-reserved range, this specification defines three representation identifiers:

- `XCDR`, which corresponds to the Extended CDR Representation encoding version 1 and takes the value 0.
- `XML`, which corresponds to the XML Data Representation and takes the value 1.
- `XCDR2`, which corresponds to Extended CDR Representation encoding version 2 and takes the value 2.

Each `Topic`, `DataReader` and `DataWriter` shall have a QoS policy `DataRepresentationQoSPolicy`. This policy shall contain a list of representation identifiers. This policy has request-offer semantics, and its value cannot be changed after the entity in question has been enabled [DDS].

- Writers offer a single representation. A writer will use its offered policy to communicate with its matched readers.

(Because the policy structure includes a sequence, it is technically possible for the writer to offer more than one representation. Implementers of this specification may use this fact in order to offer extended functionality; however, this specification does not specify any meaning for the representation identifiers after the first, and implementations may ignore them.)

- Writers belonging to implementations of `XTYPES` version 1.1 or earlier shall not announce the `XCDR2` representation identifier.
- Writers belonging to implementations of `XTYPES` version 1.2 and later:
 - Shall generate or include run-code that can serialize using version 2 encodings.
 - Optionally may generate or include run-code that can serialize using version 1 encodings. In this case, they shall offer the means to configure at run-time the encoding version used by the `DataWriter` and adjust the offered representation identifiers in the `DataRepresentationQoSPolicy` accordingly.

- Readers request one or more representations.
 - Readers requesting the XML Data Representation shall be prepared to receive either valid or merely well-formed XML documents. If a received document is well-formed but does not include any XML namespace declarations, the reader shall assume that the document could be validated using the XSD Type Representation of the corresponding sample's type if it were to include such namespace declarations.
 - Readers belonging to implementations of XTYPES version 1.1 or earlier shall not announce the XCDR2 representation identifier.
 - Shall generate or include run-time code that can deserialize version 2 encodings.
 - Shall request XCDR2 encoding.
 - Optionally may generate or include run-time code that can deserialize version 1 encodings. In this case they shall also request XCDR encoding in addition to XCDR2 encoding.
- When representations are specified in the `TopicQoS`, the first element of the sequence applies to writers of the Topic, and the whole sequence applies to readers of the Topic.
- If a writer's offered representation is contained within a reader's sequence, the offer satisfies the request and the policies are compatible. Otherwise, they are incompatible.

The default value of the `DataRepresentationQoSPolicy` shall be an empty list of preferences. An empty list of preferences shall be taken to be equivalent to a list containing the single element `XCDR`.

The `DataRepresentationQoSPolicy` shall not be changeable after its corresponding Entity has been enabled.

The rules defined in this clause result in a compatibility matrix shown in Table 59.

Table 59 – Compatibility matrix for the DataRepresentationQosPolicy

<i>DataWriter offered DataRepresentationId_t</i>	<i>DataReader requested DataRepresentationId_t</i>	<i>Encoding compatibility check</i>
<p>XCDR</p> <p>DataWriter will encode data according to version 1 encoding rules.</p> <p>Either the DataWriter is a legacy (xtypes 1.1) DataWriter or else it has been configured to use XCDR VERSION1.</p>	<p>XCDR</p> <p>DataReader is a legacy (xtypes 1.1) DataReader</p>	<p>Compatible.</p> <p>DataWriter finds its encoding among the ones understood by DataReader.</p> <p>DataReader finds its encoding among the ones understood by DataWriter.</p>
	<p>XCDR and XCDR2</p> <p>DataReader is a (xtypes 1.2) DataReader</p>	<p>Compatible.</p> <p>DataWriter finds its encoding among the ones understood by DataReader.</p> <p>DataReader finds its encoding among the ones understood by DataWriter.</p>
<p>XCDR2</p> <p>DataWriter will encode data according to version 2 encoding rules.</p> <p>DataWriter is a new (xtypes 1.2) DataWriter and it has been configured to use the version 2 encoding.</p>	<p>XCDR</p> <p>DataReader is a legacy (xtypes 1.1) DataReader</p>	<p>Not Compatible.</p> <p>DataWriter does not find its encoding among the ones understood by DataReader.</p> <p>DataReader does not find its encoding among the ones understood by DataWriter.</p>
	<p>XCDR and XCDR2</p> <p>DataReader is a new (xtypes 1.2) DataReader</p>	<p>Compatible.</p> <p>DataWriter finds its encoding among the ones understood by DataReader.</p> <p>DataReader finds its encoding among the ones understood by DataWriter.</p>

7.6.2.1.2 Use of the RTPS Encapsulation Identifier

As defined in the RTPS specification, a data encapsulation is identified by a two-byte value, the “encapsulation identifier” [RTPS]. RTPS also defines specific encapsulation identifier values corresponding to four encapsulations: big-endian CDR (CDR BE), little-endian CDR (CDR LE), big-endian parameter-list CDR (PL CDR BE), and little-endian parameter-list CDR (CDR PL LE). These encapsulations correspond to a choice of data representation and a byte-order encoding.

For the purposes of this specification, encapsulation identifiers where the first byte is in the range 0xC0 to 0xFF (both included) shall be reserved for definition by DDS implementations and shall be interpreted based on the RTPS vendor ID. The remaining values shall be reserved for the OMG⁷ for use in future specifications, including revisions of this specification.

Version 1.0 of this specification adds an additional encapsulation identifier corresponding to the XML Data Representation: XML, with the value {0x00, 0x04}. Since XML is a textual format, no byte-order differentiation is necessary.

Version 1.2 of this specification adds six additional encapsulation identifiers corresponding to PLAIN_CDR2, DELIMITED_CDR, and PL_CDR2 each with big endian or little endian encoding:

- Identifier CDR2_BE shall be used for PLAIN_CDR2 with big endian encoding
- Identifier CDR2_LE shall be used for PLAIN_CDR2 with little endian encoding
- Identifier D_CDR2_BE shall be used for DELIMITED_CDR with big endian encoding
- Identifier D_CDR2_LE shall be used for DELIMITED_CDR with little endian encoding
- Identifier PL_CDR2_BE shall be used for PL_CDR2 with big endian encoding
- Identifier PL_CDR2_LE shall be used for PL_CDR2 with little endian encoding

The encapsulation identifier field in an RTPS data sub-message shall be set such that it corresponds to the encoding version and the data representation of the outermost object whose state is represented in the message. The possible combinations are defined in Table 60.

Table 60 – RTPS encapsulation identifier

<i>Representation</i>	<i>Extensibility Kind</i>	<i>Encoding Version</i>	<i>Endianness</i>	<i>RTPS Encapsulation Identifier</i>	<i>Identifier value</i>
XCDR	FINAL	1	Big Endian	CDR_BE	{0x00, 0x00}
XCDR	FINAL	1	Little Endian	CDR_LE	{0x00, 0x01}
XCDR	APPENDABLE	1	Big Endian	CDR_BE	{0x00, 0x00}

⁷ Note that all RTPS-specified encapsulation identifier values fall within the OMG-reserved range.

XCDR	APPENDABLE	1	Little Endian	CDR_LE	{0x00, 0x01}
XCDR	MUTABLE	1	Big Endian	PL_CDR_BE	{0x00, 0x02}
XCDR	MUTABLE	1	Little Endian	PL_CDR_LE	{0x00, 0x03}
XCDR	FINAL	2	Big Endian	CDR2_BE	{0x00, 0x06}
XCDR	FINAL	2	Little Endian	CDR2_LE	{0x00, 0x07}
XCDR	APPENDABLE	2	Big Endian	D_CDR2_BE	{0x00, 0x08}
XCDR	APPENDABLE	2	Little Endian	D_CDR2_LE	{0x00, 0x09}
XCDR	MUTABLE	2	Big Endian	PL_CDR2_BE	{0x00, 0x0a}
XCDR	MUTABLE	2	Little Endian	PL_CDR_LE	{0x00, 0x0b}
XML	any	any	any	XML	{0x00, 0x04}

As defined in Sub Clause 10.2.1.2 titled “OMG CDR” of the RTPS specification, the Encapsulation Identifier is followed by a 16-bit options field. The options field is then followed by the data encoded using XCDR.

The XML encapsulation identifier is also followed by a 16-bit options field, which shall precede the data serialized using the XML data representation described in Sub Clause 7.4.4.

The RTPS specification does not define any settings for the 16-bit options field and further states that a receiver should not interpret it when it reads the options field. This DDS-XTYPES specification changes this defining the use of some bits in the options field.

Implementations of this specification shall set the lower order two bits of the 16 bit options field to a value that encodes the number of padding bytes from the end of the serialized payload to the 4-byte aligned offset that will start the next RTPS submessage. Specifically the last two bits shall be set to binary 00 if there was no padding, binary 01 if there was one byte of padding, binary 10 if there were two bytes of padding and binary 11 if there were three bytes of padding. This shall be interpreted by the receiver to determine where the serialized data ended.

For example assume structures `TypeA` and `TypeB` defined by the following IDL:

```

struct TypeA {
    short member1;
};

struct TypeB {
    short member1;
    char member2;
};

```

Furthermore assume an object O1 of type `TypeA` with value `O1.member1 = 0x11` and an object O2 of type `TypeB` with value `O2.member1 = 0x23` and `O2.member2 = 'b'`. The CDR big endian representation of these two objects, including Encapsulation header and options would be:

Object O1 representation:

```

0...2...4.....8.....16.....24.....32
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   CDR_BE { 0x00, 0x00 }           |   options { 0x00, 0x02 }           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           O1.member1 = 0x11       | padding (2 bytes) {0x00, 0x00}|
+-----+-----+-----+-----+-----+-----+-----+-----+
NEXT RTPS SUBMESSAGE...

```

Object O2 representation:

```

0...2...4.....8.....16.....24.....32
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   CDR_BE { 0x00, 0x00 }           |   options { 0x00, 0x01 }           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           O2.member1 = 0x23       | O2.member2 = 'b' | padding {0x00}|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
NEXT RTPS SUBMESSAGE...

```

7.6.2.1.3 `DataRepresentationQosPolicy`: Platform-Specific API

The conceptual model defined above shall be transformed into the IDL definitions `RepresentationId_t`, `RepresentationIdSeq`, `DATA_REPRESENTATION_QOS_POLICY_ID`, `DATA_REPRESENTATION_QOS_POLICY_NAME`, and `DataRepresentationQosPolicy`. These definitions are given in “Annex D: DDS Built-in Topic Data Types.”

The topic, publication, and subscription built-in topic data types shall each indicate the data representation of the associated entity with a new member:

```
@id(0x0073) DDS::DataRepresentationQosPolicy representation;
```

7.6.2.2 Discovery Built-in Topics

7.6.2.2.1 Type Information

A DDS `DomainParticipant` needs to have type information on remote `DomainParticipant` Topics that are also being published or subscribed by the local `DomainParticipant`. That way the `DomainParticipant` can ensure type compatibility with the remote endpoints it matches.

XTYPES 1.1 optionally included the `TypeObject` information into the Publication and Subscription discovery built-in topic data. The `TypeObject` in XTYPES version 1.1 (`TypeObjectV1`) was defined as a “library” that contained not only the data-type for the Topic-Type, but also any data-types that were recursively needed to understand the Topic-Type (e.g.

the data-types of the members of a structure). That way a DomainParticipant that discovered the endpoint would have all the type information readily available.

XTYPES 1.2 redefines the structure of the `TypeObject` (`TypeObjectV2`) and introduces a different mechanism that avoids sending `TypeObjects` to `DomainParticipants` that are not interested in it (e.g. they already know the `TypeObject`, or they are not publishing or subscribing an affected `Topic`). The XTYPES 1.2 approach is:

- Send `TypeInformation` that include `TypeIdentifiers` (instead of `TypeObjects`) in the discovery built-in topics.
- Uses the `TypeIdentifiers` to determine which types a `DomainParticipant` is interested in.
- Uses a new pair of built-in endpoints to request the `TypeObjects` for those `TypeIdentifiers` the `DomainParticipant` is interested in, and receive the reply.

The content of the type information is defined in the IDL below:

```
@extensibility(APPENDABLE)
struct TypeIdentfierWithSize {
    TypeIdentifier    type_id;
    unsigned long    typeobject_serialized_size;
};

@extensibility(APPENDABLE)
struct TypeIdentifierWithDependencies {
    TypeIdentfierWithSize    typeid_with_size;
    // The total additional types related to minimal_type
    long dependent_typeid_count;
    sequence<TypeIdentfierWithSize> dependent_typeids;
};

typedef sequence<TypeIdentifierWithDependencies>
    TypeIdentifierWithDependenciesSeq;

@extensibility(MUTABLE)
struct TypeInformation {
    @id(0x1001) TypeIdentifierWithDependencies minimal;
    @id(0x1002) TypeIdentifierWithDependencies complete;
};

typedef sequence<TypeInformation> TypeInformationSeq;
```

The `TypeInformation` includes information on the data-type associated with the `Endpoint` (`DataWriter` or `DataReader`, i.e. the `TopicType`). It includes two fields, `minimal` and `complete`.

The field *minimal* contains the MINIMAL Hash `TypeIdentifiers` for the `TopicType` and types that it depends on:

- The field *minimal.typeid_with_size* shall contain the MINIMAL Hash `TypeIdentifier` of the `TopicType` and the serialized size of the associated `TypeObject`.
- The field *minimal.dependent_typeid_count* shall contain the total number of other MINIMAL Hash `TypeIdentifiers` that correspond to data-types the `TopicType` depends on. This field may be set to -1 to indicate it is not being provided.
- The field *minimal.dependent_typeids* may contain some of the MINIMAL Hash `TypeIdentifiers` of the types the `TopicType` depends on, along with the serialized size of the respective `TypeObjects`.

The field *complete* contains the COMPLETE Hash `TypeIdentifiers` for the `TopicType` and types that it depends on:

- The field *complete.typeid_with_size* shall contain the COMPLETE Hash `TypeIdentifier` of the `TopicType` and the serialized size of the associated `TypeObject`.
- The field *complete.dependent_typeid_count* shall contain the total number of other COMPLETE Hash `TypeIdentifiers` that correspond to data-types the `TopicType` depends on. This field may be set to -1 to indicate it is not being provided.
- The field *complete.dependent_typeids* may contain some of the COMPLETE Hash `TypeIdentifiers` of the types the `TopicType` depends on, along with the serialized size of the respective `TypeObjects`.

As mentioned the field *dependent_typeids* may be used to optionally announce some of the Hash `TypeIdentifiers` the `TopicType` recursively depends on. The decision of which types to include in the *dependent_typeids* is left to the implementation: It may be set to the empty sequence, or include all the Hash `TypeIdentifiers` that the `TopicType` depends on, or something in between. If *dependent_typeid_count* is not -1, then length of the *dependent_typeids* sequence shall be less or equal to *dependent_typeid_count*.

The `TypeIdentifiers` included in the `TypeInformation` shall include only **direct HASH** `TypeIdentifiers` (see Clause 7.3.4.6.3). In addition it shall not contain individual type identifiers for types belonging to Strongly Connected Component (i.e. those with discriminator `TI_STRONG_COMPONENT`), instead it shall include the identifier of the whole Strongly-Connected Component (`SCCIdentifier`, see Clause 7.3.4.9.3).

A `DomainParticipant` can use the `TypeInformation` to determine if it already knows the associated `TopicType` and determine the type compatibility with local endpoints. In case some of the `TypeIdentifiers` announced by a remote endpoint are not known to a `DomainParticipant`, it can use the built-in `TypeLookup Service` to retrieve the `TypeObject` of the types associated with those `TypeIdentifiers`.

7.6.2.2.2 Additional members included in discovery built-in Topics

The topic, publication, and subscription built-in topic data structures shall each indicate the type(s) used for communication by the associated entity. These declarations shall be as follows:

```
@id(0x0007) ObjectName type_name;
```



```
@id(0x0072) @optional TypeObjectV1 type; // XTYPES 1.1
@id(0x0075) @optional XTypes::TypeInformation type_information; // XTYPES 1.2
```

`TypeObjectV1` corresponds to the `TypeObject` data type specified in "Annex B: Representing Types with `TypeObject`" of DDS-XTYPES Version 1.1 [DDS-XTYPES11]. Likewise, the `type` member shall be set as specified in Clause 7.3.4 of [DDS-XTYPES11].

Non-normative note: When the `TypeObjectV1` and `TypeInformation` members (called `type` and `type_information`) are omitted from the built-in topic samples, `type_name` is the only way to resolve entity matching and as a consequence, it is possible that incompatibility between topic-types is not recognized.

7.6.2.3 Built-in TypeLookup service

7.6.2.3.1 Introduction

This specification defines two built-in Topics that are used to query DomainParticipant for type information. This includes getting the `TypeObjects` associated with `TypeIdentifiers` as well as determining the list of types that a given type depends on recursively:

- One built-in topic is used for `TypeLookup` requests. It has two built-in endpoints, a `DataWriter` to send the request and a `DataReader` to receive that request.
- The second built-in topic is used for `TypeLookup` replies. It has two built-in endpoints, a `DataWriter` to send the reply and a `DataReader` to receive that reply.

The data types associated with the `TypeLookup Request/Reply` topics are defined in accordance with the Basic Service Mapping from the [DDS-RPC] specification. It is not, however, a requirement to implement the DDS-RPC specification in order to claim compliance with this specification. The only requirement is to implement the `TypeLookup` built-in endpoints as defined in this XTYPES specification.

In order to facilitate the reading of this specification, some type definitions from DDS-RPC Clause 7.5.1.1.1 have been copied in the next clause.

7.6.2.3.2 Types reused from DDS-RPC

```
/* END of definitions copied from DDS-RPC */
module dds {
typedef octet GuidPrefix_t[12];

struct EntityId_t {
    octet entityKey[3]; octet entityKind;
};

struct GUID_t {
    GuidPrefix_t guidPrefix;
    EntityId_t entityId;
};
};
```

```

};

struct SequenceNumber_t {
    long high;
    unsigned long low;
};

struct SampleIdentity {
    GUID_t writer_guid;
    SequenceNumber_t sequence_number;
};

} // module dds

// Module dds::rpc

module dds { module rpc {
    typedef octet UnknownOperation;
    typedef octet UnknownException;
    typedef octet UnusedMember;
};

enum RemoteExceptionCode_t {
    REMOTE_EX_OK,
    REMOTE_EX_UNSUPPORTED,
    REMOTE_EX_INVALID_ARGUMENT,
    REMOTE_EX_OUT_OF_RESOURCES,
    REMOTE_EX_UNKNOWN_OPERATION,
    REMOTE_EX_UNKNOWN_EXCEPTION
};

typedef string<255> InstanceName;

struct RequestHeader {
    SampleIdentity_t requestId;
    InstanceName instanceName;
};

```

```

struct ReplyHeader {
    dds::SampleIdentity relatedRequestId;
    dds::rpc::RemoteExceptionCode_t remoteEx;
};
} } // module dds::rpc
/* END of definitions copied from DDS-RPC */

```

7.6.2.3.3 TypeLookup Types and Endpoints

Compliant implementations shall include the four built-in service endpoints shown in Table 61 below.

Table 61 – Built-in Endpoints added by the XTYPES specification

<i>Built-in Endpoint</i>	<i>RTPS EntityId t</i>	<i>Associated Topic Data</i>
TypeLookupServiceRequestDataWriter	ENTITYID_TL_SVC_REQ_WRITER = {{00, 03, 00}, c3}	TypeLookup_Request
TypeLookupServiceRequestDataReader	ENTITYID_TL_SVC_REQ_READER = {{00, 03, 00}, c4}	TypeLookup_Request
TypeLookupServiceReplyDataWriter	ENTITYID_TL_SVC_REPLY_WRITER = {{00, 03, 01}, c3}	TypeLookup_Reply
TypeLookupServiceReplyDataReader	ENTITYID_TL_SVC_REPLY_READER = {{00, 03, 01}, c4}	TypeLookup_Reply

The pair `TypeLookupServiceRequestDataWriter` and `TypeLookupServiceReplyDataReader` is used to invoke the built-in `TypeLookup Service` (send the request and receive the reply).

The pair `TypeLookupServiceRequestDataReader` and `TypeLookupServiceReplyDataWriter` is used to implement the `TypeLookup Service` (receive the request and send the reply).

The Quality of Service for the four-built-in endpoints shall match the default Qos for service endpoints defined in Clause 7.10.2 of [DDS-RPC], specifically the RELIABILITY policy shall be `DDS_RELIABLE_RELIABILITY_QOS`, the HISTORY policy to `DDS_KEEP_ALL_HISTORY_QOS` and the DURABILITY policy to `DDS_VOLATILE_DURABILITY_QOS`.

The associated data-types are defined using IDL below.

```

module dds { module builtin {
const long TypeLookup_getTypes_Hash = 0xd35282d1; // @hashid("getTypes")
const long TypeLookup_getDependencies_Hash = 0x31fbaa35;
//@hashid("getDependencies");

// Query the TypeObjects associated with one or more TypeIdentifiers

```

```

@extensibility(MUTABLE)
struct TypeLookup_getTypes_In {
    @hashid sequence<TypeIdentifier>  type_ids;
};

@extensibility(MUTABLE)
struct TypeLookup_getTypes_Out {
    @hashid sequence<TypeIdentifierTypeObjectPair>  types;
    @hashid sequence<TypeIdentifierPair>  complete_to_minimal;
};

union TypeLookup_getTypes_Result switch(long) {
    case DDS_RETCODE_OK:
        TypeLookup_getTypes_Out  result;
};

// Query TypeIdentifiers that the specified types depend on
@extensibility(MUTABLE)
struct TypeLookup_getTypeDependencies_In {
    @hashid sequence<TypeIdentifier>  type_ids;
    @hashid sequence<octet, 32>      continuation_point;
};

@extensibility(MUTABLE)
struct TypeLookup_getTypeDependencies_Out {
    @hashid sequence<TypeIdentifierWithSize> dependent_typeids;
    @hashid sequence<octet, 32>          continuation_point;
};

union TypeLookup_getTypeDependencies_Result switch(long){
    case DDS_RETCODE_OK:
        TypeLookup_getTypeDependencies_Out  result;
};

// Service Request
union TypeLookup_Call switch(long) {
    case TypeLookup_getTypes_Hash:

```

```

        TypeLookup_getTypes_In          getTypes;
    case TypeLookup_getDependencies_Hash:
        TypeLookup_getTypeDependencies_In  getTypeDependencies;
};

@RPCRequestType
struct TypeLookup_Request {
    dds::rpc::RequestHeader header;
    TypeLookup_Call          data;
};

// Service Reply
union TypeLookup_Return switch(long) {
    case TypeLookup_getTypes_Hash:
        TypeLookup_getTypes_Result  getType;

    case TypeLookup_getDependencies_Hash:
        TypeLookup_getTypeDependencies_Result  getTypeDependencies;
};

@RPCReplyType
struct TypeLookup_Reply {
    dds::rpc::RequestHeader header;
    TypeLookup_Return        return;
};
}} // dds::builtin

```

The “_In” and “_Out” types are used to represent the request and reply parameters to the service. These types are defined with extensibility kind MUTABLE. Therefore they can be modified without breaking interoperability.

Implementers may add their own members to these MUTABLE types. If they do they shall use member IDs obtained using the @hashid annotation with a string value that has an Internet domain name owned by the implementor prefix. This avoids member ID conflicts with additions from other implementations. For example:

```

// Implementation from company acme.com adds parameters
// extra1 and extra2 to the getTypes request.
struct TypeLookup_getTypes_In {
    @hashid sequence<TypeIdentifier>    type_ids;
    @hashid("acme.com/extra1")    long    extra1;
}

```

```
@hashid("acme.com/extra2") string extra2;
};
```

7.6.2.3.4 Use of the TypeLookup Service

The DDS Interoperability Wire Protocol [RTPS] specifies that the `ParticipantBuiltinTopicData` shall contain the attribute called `availableBuiltinEndpoints` that is used to announce the built-in endpoints that are available in the `DomainParticipant`. See Clause 8.5.3.2 of [RTPS]. The type for this attribute is an array of `BuiltinEndpointSet_t`.

For the UDP/IP PSM the `BuiltinEndpointSet` is mapped to a bitmap represented as type `UInt32`. Each built-in endpoint is represented as a bit in this bitmap with the bit values defined in Table 9.4 (Clause 9.3.2) of [RTPS].

This DDS XTypes specification reserves additional bits to indicate the presence of the corresponding built-in end points for the `TypeObjectLookup` Service. These bits shall be set on the `availableBuiltinEndpoints`. The bit that encodes the presence of each individual endpoint is defined in Table 62 below.

Table 62 – Mapping of the built-in endpoints added by this specification to the availableBuiltinEndpoints

<i>Built-in Endpoint</i>	<i>Bit in the ParticipantBuiltinTopicData availableBuiltinEndpoints</i>
<code>TypeLookupServiceRequestDataWriter</code>	$(0x00000001 \ll 12)$
<code>TypeLookupServiceRequestDataReader</code>	$(0x00000001 \ll 13)$
<code>TypeLookupServiceReplyDataWriter</code>	$(0x00000001 \ll 14)$
<code>TypeLookupServiceReplyDataReader</code>	$(0x00000001 \ll 15)$

Participants implementing (as a server) the `TypeLookup` service shall implement the `TypeObjectServiceRequestDataReader` and `TypeObjectServiceReplyDataWriter`.

The Service *instanceName* that appears in the `dds::rpc::RequestHeader` shall be set to the string obtained by concatenating the prefix “`dds.builtin.TOS.`” With the 16-character string version of the `DomainParticipant` GUID encoded using hexadecimal digits with lower case letters. There shall be no “0x” ahead of the hexadecimal digits. For example, “`dds.builtin.TOS.123456789abcdef0`”

Participants using (as a client) the `TypeLookup` shall implement the `TypeObjectServiceRequestDataWriter` and `TypeObjectServiceReplyDataReader`.

If a participant implements the `TypeLookup` it shall respond to requests for any `TypeIdentifier` that it announced within the `TypeInformation` included in the `PublicationBuiltinTopicData` or `SubscriptionBuiltinTopicData`.

The `dds::rpc::RequestHeader` in the `TypeLookup_Request` and the `TypeLookup_Reply` shall be set as specified in the [DDS-RPC] specification.

7.6.2.3.4.1 Service operation `getTypeDependencies`

When a `DomainParticipant` receives an incomplete list of `TypeIdentifiers` in a `PublicationBuiltinTopicData` or `SubscriptionBuiltinTopicData`, it may request the additional type dependencies by invoking the `getTypeDependencies` operation.

The `TypeLookup_getTypeDependencies_In` structure shall be filled as follows:

- The field `type_ids` shall contain the sequence of `TypeIdentifiers` for which the Participant wants to get the dependencies.
 - The `TypeIdentifiers` shall be only **direct HASH** Identifiers.
 - The `TypeIdentifiers` shall be either all **MINIMAL** hash `TypeIdentifiers` or all **COMPLETE** hash `TypeIdentifiers`. That is there shall be not be mixed.
 - The `TypeIdentifiers` shall not include identifiers for individual types in **Strongly Connected Components (SCCs)**. Instead it shall use the identifier for the whole SCC (**SCCIdentifier**, see Clause 7.3.4.9.3).
- The field `continuation_point` shall not be present if the requester wants the response to include all the types that the specified types in `type_ids` depend on. Otherwise it shall be set to the `continuation_point` of the `TypeLookup_getTypeDependencies_Out` received in response to a previous call to `getTypeDependencies` with the same `type_ids`. This mechanism is used when the response of the service to a previous call to `getDependencies` did not return all the types and provided a `continuation_point`.

The `TypeLookup_getTypeDependencies_Out` structure shall be filled as follows:

- The field `dependent_typeids` shall exclusively contain of **direct HASH** `TypeIdentifiers` that are recursive dependencies from at least one of the `TypeIdentifiers` in the request.
- The field `continuation_point` shall not be present if the response contains the complete list of types, otherwise it shall contain an opaque value that the requester shall use in a subsequent request for type identifiers.

7.6.2.3.4.2 Service operation `getTypes`

A `DomainParticipant` may invoke the operation `getTypes` to retrieve the `TypeObjects` associated with a list of `TypeIdentifiers`.

A `DomainParticipant` may find out about `TypeIdentifiers` of interest as part of the information received in a `PublicationBuiltinTopicData` or `SubscriptionBuiltinTopicData`. It may also find out `TypeIdentifiers` in reply to a `getDependencies` request, or it may find them inside `TypeObjects` received in reply to a `getTypes` request. Regardless of the source it can use the `getTypes` to get the associated `TypeObjects`.

The `TypeLookup_getTypes_In` structure shall be filled as follows:

- The field `type_ids` shall contain the **direct HASH** `TypeIdentifiers` for which the participant is requesting the `TypeObjects`.
- The field `type_ids` shall not include individual `TypeIdentifiers` belonging to a Strongly Connected Component (SCC). Instead it shall use the identifier for the whole SCC (`SCCIdentifier`, see Clause 7.3.4.9.3).

The `TypeLookup_getTypes_Out` structure shall be filled as follows:

- The field `types` shall contain `TypeObjects` that correspond to the `TypeIdentifiers` in the request.
 - If the request had a **COMPLETE** `TypeIdentifiers`, the `types` shall contain **COMPLETE** `TypeObjects`.
 - If the request had **MINIMAL** `TypeIdentifiers` the `types` may contain either **MINIMAL** or **COMPLETE** `TypeObjects`.
 - The field `complete_to_minimal` shall contain the mapping from **COMPLETE** `TypeIdentifiers` to **MINIMAL** `TypeIdentifiers` for any **COMPLETE** `TypeIdentifiers` that appear within **COMPLETE** `TypeObjects` that were sent in response to a query for a **MINIMAL** `TypeIdentifier`.
 - The use of the `complete_to_minimal` field allows an implementation to only send **COMPLETE** `TypeObjects` in response to the `getTypes` request, even if the requested `TypeIdentifiers` are **MINIMAL** `TypeIdentifiers`. The combination of a **COMPLETE** `TypeObject` and the mapping of **MINIMAL** to **COMPLETE** `TypeIdentifiers` makes it possible for the receiver to reconstruct the **MINIMAL** `TypeObject`.
- If a `TypeIdentifier` was a `SCCIdentifier` (see Clause 7.3.4.9.3), then the response shall treat the `TypeObjects` within the Strongly Connected Components atomically. Either include all in the reply or none.

7.6.2.4 Type Consistency Enforcement QoS Policy

The Type Consistency Enforcement QoS Policy defines the rules for determining whether the type used to publish a given data stream is consistent with that used to subscribe to it. It applies to `DataReaderS`.

7.6.2.4.1 `TypeConsistencyEnforcementQosPolicy`: Conceptual Model

This policy defines a *type consistency kind*, which allows applications to select from among a set of predetermined policies. The following consistency kinds are specified:

- **DISALLOW_TYPE_COERCION**: The `DataWriter` and the `DataReader` must support the same data type in order for them to communicate. (This is the degree of type consistency enforcement required by the DDS specification [DDS] prior to this specification.)

- **ALLOW_TYPE_COERCION:** The `DataWriter` and the `DataReader` need not support the same data type in order for them to communicate as long as the reader's type is assignable from the writer's type.

Further details of these policies are provided in Clause 7.6.2.4.2.

This policy applies only to `DataReaders`; it does not have request-offer (RxO) semantics [DDS]. The value of this policy cannot be changed after the entity in question has been enabled.

The default enforcement kind shall be `ALLOW_TYPE_COERCION`. However, when the Service is introspecting the built-in topic data declaration of a remote `DataWriter` or `DataReader` in order to determine whether it can match with a local reader or writer, if it observes that no `TypeConsistencyEnforcementQosPolicy` value is provided (as would be the case when communicating with a Service implementation not in conformance with this specification), it shall assume a kind of `DISALLOW_TYPE_COERCION`⁸. This behavior is consistent with the type member defaulting rules defined in Clause 7.2.2.4.4.5, which state that unspecified values of enumerated types take the first value defined for their type.

This policy provides a way to control whether a type can be widened or not. A type T2 is said to widen type T1 when type T2 contains non-optional fields that are not present in T1. For example, if T2 inherits from T1 then it is said that T2 widens T1. When constructing an object O2 of the wider type T2 from an object O1 of type T1 any non-optional members in O2 not present in O1 would be set to their default values. Looking at O1 this situation is not distinguishable from the members being present in O2 and set to those same default values. In some scenarios this ambiguity may not be desirable.

Note that optional members in T2 that are not present on T1 do not make T2 “wider” than T1 according to the previous definition. This is because for optional members it is possible to tell whether that member's value was sent or not.

- The `prevent_type_widening` option controls whether type widening is allowed. If the option is set to `FALSE` (the default), type widening is permitted. If the option is set to `TRUE`, it shall cause a wider type to not be assignable to a narrower type.

This policy provides ways to ignore or enforce checking of sequence bounds, strings bounds, or member names during type assignability.

- The `ignore_sequence_bounds` option controls whether sequence bounds are taken into consideration for type assignability. If the option is set to `TRUE` (the default), sequence bounds (maximum lengths) are not considered as part of the type assignability. This means that a T2 sequence type with maximum length L2 would be assignable to a T1 sequence type with maximum length L1, even if L2 is greater than L1. If the option is set to `false`, then sequence bounds are taken into consideration for type assignability and in order for T1 to be assignable from T2 it is required that $L1 \geq L2$.

⁸ **Design rationale (non-normative):** This behavior is critical to ensure that conformant and non-conformant Service implementations reach the same conclusion regarding whether or not a `DataWriter` and a given `DataReader` are using consistent types.

- The `ignore_string_bounds` option controls whether string bounds are taken into consideration for type assignability. If the option is set to `TRUE` (the default), string bounds (maximum lengths) are not considered as part of the type assignability. This means that a `T2` string type with maximum length `L2` would be assignable to a `T1` string type with maximum length `L1`, even if `L2` is greater than `L1`. If the option is set to `false`, then string bounds are taken into consideration for type assignability and in order for `T1` to be assignable from `T2` it is required that `L1 >= L2`.
- The `ignore_member_names` option controls whether member names are taken into consideration for type assignability. If the option is set to `TRUE`, member names are considered as part of assignability in addition to member IDs (so that members with the same ID also have the same name). If the option is set to `FALSE` (the default), then member names are not ignored.

The values of `prevent_type_widening`, `ignore_sequence_bounds`, `ignore_string_bounds`, and `ignore_member_names` only apply when the type consistency kind is `ALLOW_TYPE_COERCION`, otherwise the fields are treated as though `prevent_type_widening` is set to `true` and the others are set to `false`.

This policy provides a way to declare that type information must be available in order for two endpoints to match, they cannot match solely on type names. See Sub Clause 7.6.2.4.2 for more details on how matching between a `DataWriter` and `DataReader` occurs in the presence and absence of type information.

- The `force_type_validation` option requires type information to be available in order to complete matching between a `DataWriter` and `DataReader` when set to `TRUE`, otherwise matching can occur without complete type information when set to `FALSE`. The default value is `false`.

7.6.2.4.2 Rules for Type Consistency Enforcement

Implementations of this specification shall use the type-consistency-enforcement rules defined in this clause when matching a `DataWriter` with a `DataReader`, each associated with a `Topic` of the same name. These rules are based on the data types of these entities and on the type consistency kind of the `DataReader`.

The type-consistency-enforcement rules consist of two steps.

Step 1. If both the `Publication` and the `Subscription` specify a `TypeObject`, consider it first. If the `Subscription` allows type coercion, then the `type` indicated there must be assignable from the `type` of the `Publication`, taking into account the values of `prevent_type_widening`, `ignore_sequence_bounds`, `ignore_string_bounds`, and `ignore_member_names`. If the `Subscription` does not allow type coercion, then its type must be equivalent to the type of the `Publication`.

If the `subscription` allows type coercion and the `ignore_member_names` flag is `true` in `TypeConsistencyEnforcementQoSPolicy`, assignability checking shall ignore the member names in both `Subscription` and `Publication` types. I.e., only member IDs will impact assignability.

Step 2. If either the Publication or the Subscription does not provide a `TypeObject` definition, then the type names are consulted. The Subscription and Publication `type_name` fields must match exactly, as in [DDS] prior to this specification. This step will fail if `force_type_validation` is true, regardless of the type names.

If either Step 1 or Step 2 fails, then the `Topics` associated with the `DataWriter` and `DataReader` are considered to be inconsistent: the `DataWriter` and `DataReader` shall not communicate with each other, and the Service shall trigger an `INCONSISTENT_TOPIC` status change for both the `DataReader`'s `Topic` and the `DataWriter`'s `Topic`.

If both Step 1 and Step 2 succeed, then the `Topics` are considered to be consistent, and the matching shall proceed to check other aspects of endpoint matching, such as the compatibility of the QoS, as defined by the DDS specification.

Note that the `DataWriter` and the `DataReader` can each execute the algorithm independently, having access to its own metadata as well as that of the other endpoint as communicated via DDS discovery (see Clause 7.6.3). Moreover, the algorithm is such that both sides are guaranteed to arrive at the same conclusion. That is, either both succeed or both fail.

7.6.2.4.3 TypeConsistencyEnforcementQosPolicy: Platform-Specific API

The conceptual model defined above shall be transformed into the IDL definitions

```
TypeConsistencyKind, ignore_member_names,  
TYPE_CONSISTENCY_ENFORCEMENT_QOS_POLICY_ID,  
TYPE_CONSISTENCY_ENFORCEMENT_QOS_POLICY_NAME, and  
TypeConsistencyEnforcementQosPolicy. These definitions are given in “Annex D: DDS  
Built-in Topic Data Types.”
```

The subscription built-in topic data type shall indicate the type consistency requirements of the corresponding reader:

```
@id(0x0074) DDS::TypeConsistencyEnforcementQosPolicy type_compatibility;
```

7.6.3 Local API Extensions

The following sub clauses define changes in behavior to existing operations defined by [DDS].

7.6.3.1 Operation: DomainParticipant::create_topic

As defined in [DDS], a local `Topic` object is uniquely identified by its name. In implementations conforming to this specification, that restriction shall be removed. The Service may instantiate multiple objects of the same name, provided that all of them represent type-based subsets of “the same” network topic; therefore, they must have consistent QoS with one another.

7.6.3.2 Operation: DomainParticipant::lookup_topicdescription

As defined in [DDS], a local `TopicDescription` object is uniquely identified by its name. In implementations conforming to this specification, that restriction shall be removed. The definition of `lookup_topicdescription` operation shall be modified from the one in the [DDS] specification as follows.

The `lookup_topicdescription` operation shall accept an optional *in* unsigned long argument called `index`. This shall be the last argument.

When the operation is called with only `topic_name`. It shall behave as if called with `index = 0`.

When the operation is called with both a `topic_name` and an `index`, the operation shall return one of the `TopicDescription` associated with the `DomainParticipant` with a matching `topic_name`. The value of the `index` parameter shall be treated as an “iterator” over the sequence of `TopicDescription` instances that match that `topic_name`. Each value of the `index` shall return a unique (different) `TopicDescription`. Values of the `index` from 0 to one less than the number of different `TopicDescriptions` match the `topic_name` shall return a `TopicDescription` and values of the `index` outside the range shall return *nil*.

7.6.4 Built-in Types

DDS shall provide a few types preregistered “out of the box” to allow users to address certain simple use cases without the need for code generation, dynamic type definition, or type registration. These types are:

- **DDS::String**: A single unbounded string; a data type without a key.
- **DDS::KeyedString**: A pair of unbounded strings, one representing the payload and a second representing its key.
- **DDS::Bytes**: An unbounded sequence of bytes, useful for transmitting opaque or application-serialized data.
- **DDS::KeyedBytes**: A payload consisting of an unbounded sequence of bytes plus a key field, an unbounded string.

The built-in types shall be defined as in the following sections and shall be automatically registered by the Service under their fully qualified physical names (as above) with each `DomainParticipant` at the time it is enabled.

Like all non-nested types used with DDS, the built-in types shall have corresponding type-specific `DataWriter` and `DataReader` classes. These shall instantiate the type-specific operations defined by the DDS specification as defined in the following sections; they shall also provide additional overloads.

The built-in types are described briefly below; their complete definitions may be found in “Annex E: Built-in Types.”

7.6.4.1 String

The `DDS::String` type is a simple structure wrapper around a single unbounded string. The wrapper structure exists in order to provide the Service implementation with a non-nested type definition and as a basis of the `TypeObject` object propagated with the built-in topics. But the `StringDataWriter` and `StringDataReader` APIs are defined based on the built-in `string` type for convenience.

7.6.4.2 KeyedString

The `DDS::KeyedString` type is similar to `DDS::String`, but it is a keyed type; the key is an additional unbounded string. `DDS::KeyedStringDataWriter` provides additional overloads that “unwrap” this structure, allowing applications to pass the two strings directly.

7.6.4.3 Bytes

The `DDS::Bytes` type is a simple structure wrapper around a single unbounded sequence of bytes. The wrapper structure exists in order to provide the Service implementation with a non-nested type definition and as a basis of the `TypeObject` object propagated with the built-in topics. The `BytesDataWriter` API is defined based on the underlying sequence for convenience; the `BytesDataReader` API is based on `DDS::Bytes` because of the awkwardness of sequences of sequences.

7.6.4.4 KeyedBytes

The `DDS::KeyedBytes` type is similar to `DDS::Bytes`, but it is a keyed type; the key is an unbounded string. `DDS::KeyedBytesDataWriter` provides additional overloads that “unwrap” this structure, allowing applications to pass the string and sequence directly.

7.6.5 Use of Dynamic Data and Dynamic Type

Using the `DynamicData` and `DynamicType` APIs applications can publish and subscribe data of any type without having compile-time knowledge of the type.

The API is still strongly typed; each specific `Type` must be registered with the `DomainParticipant`. The `DynamicType` interface can be used to construct the `Type` and register it with the `DomainParticipant`. The `DynamicData` interface can be used to create objects of a specified `Type` (expressed by means of a `DynamicType`) and publish and subscribe data objects of that type.

In order to for an application to use a type for publication or subscription the type must first be registered with the corresponding `DomainParticipant` in the same manner as a type defined at compile time.

7.6.5.1 Type Support

Application code (i.e. business logic) generally depends statically on particular types and their members. In contrast, infrastructure code (i.e. logic that is independent of particular applications) generally must not depend on application-specific types, because such dependencies prevent that code from being reused. These two kinds of code can exist within a single component.

Therefore, it is desirable to allow conversions among static and dynamic bindings for the same types and samples. These conversions shall be provided by operations on the generic `TypeSupport` interface and its extended interfaces.

7.6.5.1.1 TypeSupport Interface

The following operations shall be added to the `TypeSupport` interface defined by [DDS]. (The operations on this interface already defined in [DDS] are unchanged.)

Table 63 – New TypeSupport operations

Operations		
get_type		DynamicType

7.6.5.1.1.1 Operation: get_type

Get a `DynamicType` object corresponding to the `TypeSupport`'s data type.

7.6.5.1.2 FooTypeSupport Interface

The following operations shall be added to the `FooTypeSupport` interface defined by [DDS]. (The operations on this interface already defined in [DDS] are unchanged.)

Table 64 – New FooTypeSupport operations

Operations		
create_sample		Foo
	src	DynamicData
create_dynamic_sample		DynamicData
	src	Foo

7.6.5.1.2.1 Operation: create_sample

Create a sample of the `TypeSupport`'s data type with the contents of an input `DynamicData` object.

Parameter `src` – The source object whose contents are to be reflected in the resulting object. This method shall fail with a nil return result if this object is nil or if the `DynamicType` of this object is not compatible with the `TypeSupport`'s data type.

7.6.5.1.2.2 Operation: create_dynamic_sample

Create a `DynamicData` object with the contents of an input sample of the `TypeSupport`'s data type.

Parameter `src` – The source object whose contents are to be reflected in the resulting object. This method shall fail with a nil return result if this object is nil.

7.6.5.1.3 DynamicTypeSupport

The `DynamicTypeSupport` interface extends the `FooTypeSupport` interface defined by the DDS specification where “Foo” is the type `DynamicData`.

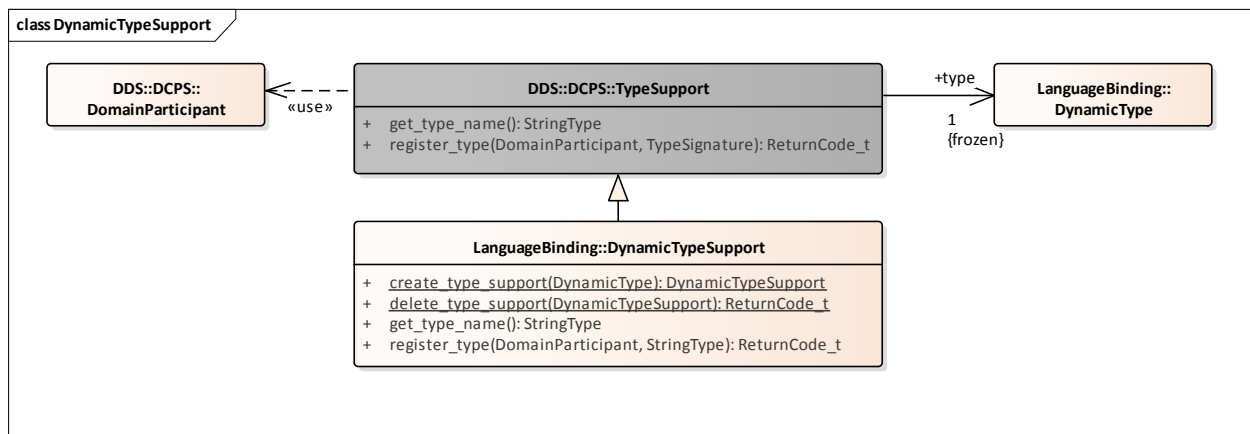


Figure 32 – Dynamic Type Support

Table 65 – DynamicTypeSupport properties and operations

<i>DynamicTypeSupport</i>		
Operations		
register_type		ReturnCode_t
	participant	DomainParticipant
	type name	string<Char8,256>
get_type_name		string<Char8,256>
static create_type_support		DynamicTypeSupport
	type	DynamicType
static delete_type_support		ReturnCode_t
	support	DynamicTypeSupport

7.6.5.1.4 Operations: register_type, get_type_name

These operations are defined by, and described in, the DDS specification.

7.6.5.1.5 Operation: create_type_support

Create and return a new `DynamicTypeSupport` object capable of registering the given type with DDS `DomainParticipant`s. The implementation shall ensure that the new type support has a “copy” of the given type object, such that subsequent changes to, or deletions of, the argument object do not impact the new type support. All objects returned by this operation should eventually be deleted by calling `delete_type_support`.

If an error occurs, this method shall return a nil value.

Parameter type - The type for which to create a type support. If this argument is nil or is a nested type, the operation shall fail and return a nil value.

7.6.5.1.6 Operation: `delete_type_support`

Delete the given type support object, which was previously created by this factory.

If this argument is nil, the operation shall return successfully without having any observable effect.

Parameter `type_support` – The type support object to delete. If this argument is an object that was already deleted, and the implementation is able to detect that fact (which is not required), this operation shall fail with `RETCODE_ALREADY_DELETED`. If an implementation-specific error occurs, this method shall fail with `RETCODE_ERROR`.

7.6.5.2 `DynamicDataWriter` and `DynamicDataReader`

The `DynamicDataWriter` interface instantiates the `FooDataWriter` interface defined by the DDS specification where “Foo” is the type `DynamicData`.

The `DynamicDataReader` interface instantiates the `FooDataReader` interface defined by the DDS specification where “Foo” is the type `DynamicData`.

These types do not define additional properties or operations.

7.6.6 DCPS Queries and Filters

[DDS] defines the syntax for content-based filters, queries, and joins in “*Annex A: Syntax for DCPS Queries and Filters*”. This syntax shall be extended as follows.

7.6.6.1 Member Names

[DDS] Clause A.2 defines the syntax for referring to a member of a (potentially nested) data structure. Such a reference is known as a `FIELDNAME`. The syntax shall be extended as follows:

- *Arrays and sequences*: Elements in these ordered collections shall be indicated by a zero-based subscript enclosed in square brackets, e.g. `my_collection[0]`. Such an expression shall be considered to have the type that is the element type of the collection.
- *Maps*: Value elements in these unordered collections shall be indicated by a string representation of a corresponding key element, according to the syntax of `STRING`, enclosed in square brackets, e.g. `my_map['key']`. The key shall be expressed as a string even if the map’s key type is an integer type; this distinguishes a map lookup from an index into an ordered collection. Such an expression shall be considered to have the type that is the value element type of the map.
- *Bitmasks*: A flag in a bitmask shall be indicated by its name, according to the syntax of `ENUMERATEDVALUE`, enclosed in square brackets, e.g. `my_bitmask['MY_FLAG']`. Such an expression shall be considered to have a Boolean type: true if the bit is set or false if it is not. Comparisons with the integer literals 1 and 0 shall also be allowed.

7.6.6.2 Optional Type Members

A member of an aggregated type may be compared to the special value `null`. Such comparisons obey the following rules:

- If the member is optional, and it takes no value in the given object, it shall be considered equal to `null`.
- If the member is optional, and it does take a value in the given object, it shall not be considered equal to `null`.
- No non-optional member shall ever be considered equal to `null`.

Inequalities expressed relative to `null` shall never evaluate to true—no value is greater than or less than `null`.

7.6.6.3 Grammar Extensions

The `Parameter` production in the grammar given in [DDS] Clause A.1 shall be redefined as follows:

```
Parameter ::=
    | CHARVALUE
    | FLOATVALUE
    | STRING
    | ENUMERATEDVALUE
    | BOOLEANVALUE
    | NULLVALUE
    | PARAMETER
    .
```

(New tokens have been highlighted in **bold**.)

The `BOOLEANVALUE` token shall be either `true` or `false` (case-insensitive).

The `NULLVALUE` token shall always be `null`.

7.6.7 Interoperability of Keyed Topics

As described in [RTPS] Clause 9.6.3.3, “KeyHash (PID_KEY_HASH)”, the key hash for a given object of a keyed type is obtained by first serializing the values of the key members in their declaration order. The algorithm described in that clause shall be amended such that key member values shall be serialized in the ascending orders of their member IDs. For calculation of KeyHash for mutable types, the key members shall be serialized without any parameter encapsulation.

Design rationale (non-normative): This change ensures that key hash values remain stable in the face of member order permutations. It is backwards compatible, because this specification

interprets all pre-existing type definitions (which lack explicit member IDs) as implying member IDs in declaration order. Thus all pre-existing key hashing algorithm implementations already conform to this specification when applied to pre-existing type definitions. Further, ignoring parameter encapsulation for mutable types avoids ambiguities with respect to using short/long parameter encapsulation. For mutable types, the key members are serialized as if the top-level and nested types were declared appendable.

8. Changes or Extensions Required to Adopted OMG Specifications

8.1 Extensions

8.1.1 DDS

This specification extends the DDS specification [DDS] as described in Clause 2.1, “Programming Interface Conformance,” above. As described in that clause, these extensions comprise a new, optional conformance level within the DDS specification.

This specification *does not* modify or invalidate any pre-existing DDS profiles or conformance levels, including the Minimum Profile. Therefore, previously conformant DDS implementations remain conformant, and conformance to this additional specification by DDS implementations is completely optional.

8.2 Changes

This specification does not change any pre-existing programming interface, behavior, or other facility of any adopted OMG specification.

Annex A: XML Type Representation Schema

The following set of XML Schema Documents (XSD) formally defines the structure of XML documents conforming to the XML Type Representation.

The first schema file, `dds_types.xsd`, declares the appropriate `targetNamespace` for this specification (i.e., `http://www.omg.org/dds`), includes a schema containing the types definition called `dds_types_definition.xsd`, and defines the root element for XML documents containing type definitions.

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- dds_xtypes.xsd -->
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns="http://www.omg.org/dds"
  targetNamespace="http://www.omg.org/dds"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">
  <xs:include schemaLocation="dds_types_definition.xsd" />
  <xs:element name="types" type="typeLibrary"/>
</xs:schema>
```

The types definition schema file does not declare a `targetNamespace`, which makes it simpler for other specifications to include the schema file without having to deal with namespace declarations. This follows the so-called Chameleon Namespace Design, in which the schema with no `targetNameSpace` takes the "color" (namely, the `targetNamespace`) of the XSD file that includes it.

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- dds_types_definition.xsd -->
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">
  <!-- ===== -->
  <!-- Identifiers -->
  <!-- ===== -->

  <xs:simpleType name="identifierName">
    <xs:restriction base="xs:string">
      <xs:pattern value="([a-zA-Z]|:|:)([a-zA-Z_0-9]|:|:)*"/>
    </xs:restriction>
  </xs:simpleType>
```

```

<!-- ===== -->
<!-- File Inclusion -->
<!-- ===== -->

<xs:simpleType name="fileName">
  <xs:restriction base="xs:string">
    </xs:restriction>
  </xs:simpleType>

<xs:complexType name="includeDecl">
  <xs:attribute name="file"
    type="fileName"
    use="required"/>
</xs:complexType>

<!-- ===== -->
<!-- Forward Declarations -->
<!-- ===== -->

<xs:simpleType name="forwardDeclTypeKind">
  <xs:restriction base="xs:string">
    <xs:enumeration value="enum"/>
    <xs:enumeration value="struct"/>
    <xs:enumeration value="union"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="forwardDecl">
  <xs:attribute name="name"
    type="identifierName"
    use="required"/>
  <xs:attribute name="kind"
    type="forwardDeclTypeKind"
    use="required"/>
</xs:complexType>

<!-- ===== -->

```

```

<!-- Basic Types -->
<!-- ===== -->

<xs:simpleType name="allTypeKind">
  <xs:restriction base="xs:string">
    <!-- Primitive Types -->
    <xs:enumeration value="boolean"/>
    <xs:enumeration value="byte"/>
    <xs:enumeration value="char8"/>
    <xs:enumeration value="char16"/>
    <xs:enumeration value="int16"/>
    <xs:enumeration value="uint16"/>
    <xs:enumeration value="int32"/>
    <xs:enumeration value="uint32"/>
    <xs:enumeration value="int64"/>
    <xs:enumeration value="uint64"/>
    <xs:enumeration value="float32"/>
    <xs:enumeration value="float64"/>
    <xs:enumeration value="float128"/>

    <!-- String containers -->
    <xs:enumeration value="string"/>
    <xs:enumeration value="wstring"/>

    <!-- Some other type -->
    <xs:enumeration value="nonBasic"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="arrayDimensionsKind">
  <xs:restriction base="xs:string">
  </xs:restriction>
</xs:simpleType>

<!-- ===== -->
<!-- Constants -->

```

```
<!-- ===== -->
```

```
<xs:complexType name="constDecl">  
  <xs:attribute name="name"  
    type="identifierName"  
    use="required"/>  
  <xs:attribute name="type"  
    type="allTypeKind"  
    use="required"/>  
  <xs:attribute name="nonBasicTypeName"  
    type="identifierName"  
    use="optional"/>  
  <xs:attribute name="value"  
    type="xs:string"  
    use="required"/>  
</xs:complexType>
```

```
<!-- ===== -->
```

```
<!-- Aggregated Types (General) -->
```

```
<!-- ===== -->
```

```
<xs:simpleType name="memberId">  
  <xs:restriction base="xs:unsignedInt">  
    <xs:minInclusive value="0"/>  
    <xs:maxInclusive value="268435455"/><!-- 0x0FFFFFFF -->  
  </xs:restriction>  
</xs:simpleType>
```

```
<xs:complexType name="simpleMemberDecl">  
  <xs:attribute name="name"  
    type="identifierName"  
    use="required"/>  
  
  <xs:attribute name="type"  
    type="allTypeKind"  
    use="required"/>  
  
  <xs:attribute name="nonBasicTypeName"
```



```

        type="identifierName"
        use="optional"/>
</xs:complexType>

<xs:simpleType name="tryConstructKind">
  <xs:restriction base="xs:string">
    <xs:enumeration value="discard"/>
    <xs:enumeration value="use_default"/>
    <xs:enumeration value="trim"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="memberDecl">
  <xs:complexContent>
    <xs:extension base="simpleMemberDecl">
      <xs:sequence>
        <xs:element name="annotate"
          type="annotationDecl"
          minOccurs="0"
          maxOccurs="unbounded"/>
      </xs:sequence>

      <xs:attribute name="external"
        type="xs:boolean"
        use="optional"
        default="true"/>
      <xs:attribute name="tryConstruct"
        type="tryConstructKind"
        use="optional"
        default="use_default"/>
      <xs:attribute name="mapKeyType"
        type="allTypeKind"
        use="optional"/>
      <xs:attribute name="mapKeyNonBasicTypeName"
        type="identifierName"
        use="optional"/>
      <xs:attribute name="stringMaxLength"

```

```

        type="xs:string"
        use="optional"/>
<xs:attribute name="mapKeyStringMaxLength"
        type="xs:string"
        use="optional"/>
<xs:attribute name="sequenceMaxLength"
        type="xs:string"
        use="optional"/>
<xs:attribute name="mapMaxLength"
        type="xs:string"
        use="optional"/>
<xs:attribute name="arrayDimensions"
        type="arrayDimensionsKind"
        use="optional"/>

</xs:extension>
</xs:complexContent>
</xs:complexType>

<xs:complexType name="verbatimDecl">
  <xs:sequence>
    <xs:element name="text"
      type="xs:string"
      minOccurs="1"
      maxOccurs="1"/>
  </xs:sequence>

  <xs:attribute name="language"
    type="xs:string"
    use="optional"
    default="*" />
  <xs:attribute name="placement"
    type="xs:string"
    use="optional"
    default="before-declaration" />
</xs:complexType>

```

```

<xs:simpleType name="extensibilityKind">
  <xs:restriction base="xs:string">
    <xs:enumeration value="final"/>
    <xs:enumeration value="appendable"/>
    <xs:enumeration value="mutable"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="autoIdKind">
  <xs:restriction base="xs:string">
    <xs:enumeration value="hash"/>
    <xs:enumeration value="sequential"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="structOrUnionTypeDecl">
  <xs:sequence>
    <xs:choice minOccurs="0" maxOccurs="unbounded">
      <xs:element name="annotate"
        type="annotationDecl"/>
      <xs:element name="verbatim"
        type="verbatimDecl"/>
    </xs:choice>
  </xs:sequence>

  <xs:attribute name="name"
    type="identifierName"
    use="required"/>
  <xs:attribute name="nested"
    type="xs:boolean"
    use="optional"
    default="false"/>
  <xs:attribute name="extensibility"
    type="extensibilityKind"
    use="optional"
    default="appendable"/>
  <xs:attribute name="autoid"

```

```

        type="autoIdKind"
        use="optional"
        default="hash"/>
</xs:complexType>

<!-- ===== -->
<!-- Annotations -->
<!-- ===== -->

<xs:complexType name="annotationTypeDecl">
  <xs:sequence>
    <xs:element name="member"
      type="simpleMemberDecl"
      minOccurs="0"
      maxOccurs="unbounded"/>
  </xs:sequence>

  <xs:attribute name="name"
    type="identifierName"
    use="required"/>
  <xs:attribute name="baseType"
    type="identifierName"
    use="optional"/>
</xs:complexType>

<xs:complexType name="annotationMemberValueDecl">
  <xs:attribute name="name"
    type="identifierName"
    use="required"/>
  <xs:attribute name="value"
    type="xs:string"
    use="optional"/>
</xs:complexType>

<xs:complexType name="annotationDecl">

```

```

<xs:sequence>
  <xs:element name="member"
    type="annotationMemberValueDecl"
    minOccurs="0"
    maxOccurs="unbounded"/>
</xs:sequence>

<xs:attribute name="name"
  type="identifierName"
  use="required"/>
</xs:complexType>

<!-- ===== -->
<!-- Structures -->
<!-- ===== -->

<xs:complexType name="structMemberDecl">
  <xs:complexContent>
    <xs:extension base="memberDecl">
      <xs:attribute name="id"
        type="memberId"
        use="optional"/>

      <xs:attribute name="optional"
        type="xs:boolean"
        use="optional"
        default="true"/>

      <xs:attribute name="mustUnderstand"
        type="xs:boolean"
        use="optional"
        default="true"/>

      <xs:attribute name="nonSerialized"
        type="xs:boolean"
        use="optional"
        default="true"/>

      <xs:attribute name="key"
        type="xs:boolean"

```

```

        use="optional"
        default="true"/>
    </xs:extension>
</xs:complexContent>
</xs:complexType>

<xs:complexType name="structDecl">
    <xs:complexContent>
        <xs:extension base="structOrUnionTypeDecl">
            <xs:sequence>
                <xs:choice maxOccurs="unbounded">
                    <xs:element name="member"
                                type="structMemberDecl"
                                minOccurs="1"/>
                    <xs:element name="const"
                                type="constDecl"
                                minOccurs="0"/>
                </xs:choice>
            </xs:sequence>

            <xs:attribute name="baseType"
                        type="identifierName"
                        use="optional"/>
        </xs:extension>
    </xs:complexContent>
</xs:complexType>

<!-- ===== -->
<!-- Unions -->
<!-- ===== -->

<xs:complexType name="unionMemberDecl">
    <xs:complexContent>
        <xs:extension base="memberDecl"/>
    </xs:complexContent>
</xs:complexType>

```

```

<xs:complexType name="discriminatorDecl">
  <xs:sequence>
    <xs:element name="annotate"
      type="annotationDecl"
      minOccurs="0"
      maxOccurs="unbounded"/>
  </xs:sequence>

  <xs:attribute name="type"
    type="identifierName"
    use="required"/>
  <xs:attribute name="nonBasicTypeName"
    type="identifierName"
    use="optional"/>
  <xs:attribute name="key"
    type="xs:boolean"
    use="optional"
    default="false"/>
</xs:complexType>

<xs:complexType name="caseDiscriminatorDecl">
  <xs:attribute name="value"
    type="xs:string"
    use="required"/>
</xs:complexType>

<xs:complexType name="caseDecl">
  <xs:sequence>
    <xs:element name="caseDiscriminator"
      type="caseDiscriminatorDecl"
      minOccurs="1"
      maxOccurs="unbounded"/>
    <xs:element name="member"
      type="unionMemberDecl"
      minOccurs="1"
      maxOccurs="1"/>
  </xs:sequence>
</xs:complexType>

```

```

    </xs:sequence>
</xs:complexType>

<xs:complexType name="unionDecl">
  <xs:complexContent>
    <xs:extension base="structOrUnionTypeDecl">
      <xs:sequence>
        <xs:element name="discriminator"
          type="discriminatorDecl"
          minOccurs="1"
          maxOccurs="1"/>
        <xs:element name="case"
          type="caseDecl"
          minOccurs="1"
          maxOccurs="unbounded"/>
      </xs:sequence>
    </xs:extension>
  </xs:complexContent>
</xs:complexType>

<!-- ===== -->
<!-- Aliases -->
<!-- ===== -->

<xs:complexType name="typedefDecl">
  <xs:attribute name="name"
    type="identifierName"
    use="required"/>

  <xs:attribute name="type"
    type="allTypeKind"
    use="required"/>

  <xs:attribute name="nonBasicTypeName"
    type="identifierName"
    use="optional"/>

```



```

<xs:attribute name="mapKeyType"
              type="allTypeKind"
              use="optional"/>

<xs:attribute name="mapKeyNonBasicTypeName"
              type="identifierName"
              use="optional"/>

<xs:attribute name="stringMaxLength"
              type="xs:string"
              use="optional"/>

<xs:attribute name="mapKeyStringMaxLength"
              type="xs:string"
              use="optional"/>

<xs:attribute name="sequenceMaxLength"
              type="xs:string"
              use="optional"/>

<xs:attribute name="mapMaxLength"
              type="xs:string"
              use="optional"/>

<xs:attribute name="arrayDimensions"
              type="arrayDimensionsKind"
              use="optional"/>

<xs:attribute name="external"
              type="xs:boolean"
              use="optional"/>
</xs:complexType>

<!-- ===== -->
<!-- Enumerations -->
<!-- ===== -->

```

```

<xs:simpleType name="enumBitBound">
  <xs:restriction base="xs:unsignedShort">
    <xs:minInclusive value="1"/>
    <xs:maxInclusive value="32"/>
  </xs:restriction>
</xs:simpleType>

```

```

<xs:complexType name="enumeratorDecl">
  <xs:sequence>
    <xs:element name="annotate"
      type="annotationDecl"
      minOccurs="0"
      maxOccurs="unbounded"/>
  </xs:sequence>

```

```

  <xs:attribute name="name"
    type="identifierName"
    use="required"/>

```

```

  <xs:attribute name="value"
    type="xs:string"
    use="optional"/>

```

```

  <xs:attribute name="defaultLiteral"
    type="xs:boolean"
    use="optional"
    default="true"/>

```

```

</xs:complexType>

```

```

<xs:complexType name="enumDecl">
  <xs:sequence>
    <xs:element name="annotate"
      type="annotationDecl"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xs:element name="verbatim"

```

```

        type="verbatimDecl"
        minOccurs="0"
        maxOccurs="unbounded"/>
    <xs:element name="enumerator"
        type="enumeratorDecl"
        minOccurs="1"
        maxOccurs="unbounded"/>
</xs:sequence>

<xs:attribute name="name"
    type="identifierName"
    use="required"/>
<xs:attribute name="bitBound"
    type="enumBitBound"
    use="optional"
    default="32"/>
</xs:complexType>

<!-- ===== -->
<!-- Bit Masks -->
<!-- ===== -->

<xs:simpleType name="bitmaskBitBound">
    <xs:restriction base="xs:unsignedShort">
        <xs:minInclusive value="1"/>
        <xs:maxInclusive value="64"/>
    </xs:restriction>
</xs:simpleType>

<xs:simpleType name="flagIndex">
    <xs:restriction base="xs:unsignedShort">
        <xs:minInclusive value="0"/>
        <xs:maxInclusive value="63"/>
    </xs:restriction>
</xs:simpleType>

<xs:complexType name="flagDecl">

```

```

<xs:sequence>
  <xs:element name="annotate"
    type="annotationDecl"
    minOccurs="0"
    maxOccurs="unbounded"/>
</xs:sequence>

<xs:attribute name="name"
  type="identifierName"
  use="required"/>

<xs:attribute name="position"
  type="flagIndex"
  use="required"/>
</xs:complexType>

<xs:complexType name="bitmaskDecl">
  <xs:sequence>
    <xs:element name="annotate"
      type="annotationDecl"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xs:element name="flag"
      type="flagDecl"
      minOccurs="0"
      maxOccurs="64"/>
  </xs:sequence>

  <xs:attribute name="name"
    type="identifierName"
    use="required"/>

  <xs:attribute name="bitBound"
    type="bitmaskBitBound"
    use="optional"
    default="32"/>
</xs:complexType>

```

```

<!-- ===== -->
<!-- Modules -->
<!-- ===== -->

<xs:group name="moduleElements">
  <xs:sequence>
    <xs:choice maxOccurs="unbounded">
      <xs:element name="include2"
        type="includeDecl"
        minOccurs="0"/>
      <xs:element name="forward_dcl"
        type="forwardDecl"
        minOccurs="0"/>
      <xs:element name="const"
        type="constDecl"
        minOccurs="0"/>
      <xs:element name="module"
        type="moduleDecl"
        minOccurs="0"/>
      <xs:element name="struct"
        type="structDecl"
        minOccurs="0"/>
      <xs:element name="union"
        type="unionDecl"
        minOccurs="0"/>
      <xs:element name="annotation"
        type="annotationTypeDecl"
        minOccurs="0"/>
      <xs:element name="typedef"
        type="typedefDecl"
        minOccurs="0"/>
      <xs:element name="enum"
        type="enumDecl"
        minOccurs="0"/>
      <xs:element name="bitmask"
        type="bitmaskDecl"

```

```

        minOccurs="0"/>
    </xs:choice>
</xs:sequence>
</xs:group>

<xs:complexType name="moduleDecl">
    <xs:sequence>
        <xs:element name="include"
            type="includeDecl"
            minOccurs="0"
            maxOccurs="unbounded"/>
        <xs:group ref="moduleElements"
            minOccurs="0"
            maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="name"
        type="identifierName"
        use="required"/>
    <xs:attribute name="autoid"
        type="autoIdKind"
        use="optional"
        default="hash"/>
</xs:complexType>

<xs:complexType name="typeLibrary">
    <xs:group ref="moduleElements"/>
</xs:complexType>

</xs:schema>

```

Annex B: Representing Types with TypeObject

The following IDL formally describes the `TypeObject` type and those nested types on which it depends.

```
/* dds-xtypes_typeobject.idl */

// The types in this file shall be serialized with XCDR encoding version 2
module DDS { module XTypes {

    // ----- Equivalence Kinds -----
    typedef octet EquivalenceKind;
    const octet EK_MINIMAL    = 0xF1; // 0x1111 0001
    const octet EK_COMPLETE  = 0xF2; // 0x1111 0010
    const octet EK_BOTH      = 0xF3; // 0x1111 0011

    // ----- TypeKinds (begin) -----
    typedef octet TypeKind;

    // Primitive TKs
    const octet TK_NONE      = 0x00;
    const octet TK_BOOLEAN   = 0x01;
    const octet TK_BYTE      = 0x02;
    const octet TK_INT16     = 0x03;
    const octet TK_INT32     = 0x04;
    const octet TK_INT64     = 0x05;
    const octet TK_UINT16    = 0x06;
    const octet TK_UINT32    = 0x07;
    const octet TK_UINT64    = 0x08;
    const octet TK_FLOAT32   = 0x09;
    const octet TK_FLOAT64   = 0x0A;
    const octet TK_FLOAT128  = 0x0B;
    const octet TK_CHAR8     = 0x10;
    const octet TK_CHAR16    = 0x11;

    // String TKs
    const octet TK_STRING8   = 0x20;
```

```

const octet TK_STRING16    = 0x21;

// Constructed/Named types
const octet TK_ALIAS       = 0x30;

// Enumerated TKs
const octet TK_ENUM        = 0x40;
const octet TK_BITMASK     = 0x41;

// Structured TKs
const octet TK_ANNOTATION  = 0x50;
const octet TK_STRUCTURE   = 0x51;
const octet TK_UNION       = 0x52;
const octet TK_BITSET      = 0x53;

// Collection TKs
const octet TK_SEQUENCE    = 0x60;
const octet TK_ARRAY       = 0x61;
const octet TK_MAP         = 0x62;
// ----- TypeKinds (end) -----

// ----- Extra TypeIdentifiers (begin) -----
typedef octet TypeIdentiferKind;
const octet TI_STRING8_SMALL      = 0x70;
const octet TI_STRING8_LARGE      = 0x71;
const octet TI_STRING16_SMALL     = 0x72;
const octet TI_STRING16_LARGE     = 0x73;

const octet TI_PLAIN_SEQUENCE_SMALL = 0x80;
const octet TI_PLAIN_SEQUENCE_LARGE = 0x81;

const octet TI_PLAIN_ARRAY_SMALL   = 0x90;
const octet TI_PLAIN_ARRAY_LARGE   = 0x91;

const octet TI_PLAIN_MAP_SMALL     = 0xA0;
const octet TI_PLAIN_MAP_LARGE     = 0xA1;

```



```

const octet TI_STRONGLY_CONNECTED_COMPONENT = 0xB0;
// ----- Extra TypeIdentifiers (end) -----

// The name of some element (e.g. type, type member, module)
// Valid characters are alphanumeric plus the "_" cannot start with digit
const long MEMBER_NAME_MAX_LENGTH = 256;
typedef string<MEMBER_NAME_MAX_LENGTH> MemberName;

// Qualified type name includes the name of containing modules
// using "::" as separator. No leading "::". E.g. "MyModule::MyType"
const long TYPE_NAME_MAX_LENGTH = 256;
typedef string<TYPE_NAME_MAX_LENGTH> QualifiedTypeName;

// Every type has an ID. Those of the primitive types are pre-defined.
typedef octet PrimitiveTypeId;

// First 14 bytes of MD5 of the serialized TypeObject using XCDR
// version 2 with Little Endian encoding
typedef octet EquivalenceHash[14];

// First 4 bytes of MD5 of of a member name converted to bytes
// using UTF-8 encoding and without a 'nul' terminator.
// Example: the member name "color" has NameHash {0x70, 0xDD, 0xA5, 0xDF}
typedef octet NameHash[4];

// Long Bound of a collection type
typedef unsigned long LBound;
typedef sequence<LBound> LBoundSeq;
const LBound INVALID_LBOUND = 0;

// Short Bound of a collection type
typedef octet SBound;
typedef sequence<SBound> SBoundSeq;
const SBound INVALID_SBOUND = 0;

@extensibility(FINAL) @nested
union TypeObjectHashId switch (octet) {

```

```

    case EK_COMPLETE:
    case EK_MINIMAL:
        EquivalenceHash hash;
};

// Flags that apply to struct/union/collection/enum/bitmask/bitset
// members/elements and DO affect type assignability
// Depending on the flag it may not apply to members of all types
// When not all, the applicable member types are listed
@bit_bound(16)
bitmask MemberFlag {
    @position(0) TRY_CONSTRUCT1,      // T1 | 00 = INVALID, 01 = DISCARD
    @position(1) TRY_CONSTRUCT2,      // T2 | 10 = USE_DEFAULT, 11 = TRIM
    @position(2) IS_EXTERNAL,         // X StructMember, UnionMember,
                                     // CollectionElement
    @position(3) IS_OPTIONAL,         // O StructMember
    @position(4) IS_MUST_UNDERSTAND, // M StructMember
    @position(5) IS_KEY,             // K StructMember, UnionDiscriminator
    @position(6) IS_DEFAULT          // D UnionMember, EnumerationLiteral
};

typedef MemberFlag  CollectionElementFlag; // T1, T2, X
typedef MemberFlag  StructMemberFlag;     // T1, T2, O, M, K, X
typedef MemberFlag  UnionMemberFlag;      // T1, T2, D, X
typedef MemberFlag  UnionDiscriminatorFlag; // T1, T2, K
typedef MemberFlag  EnumeratedLiteralFlag; // D
typedef MemberFlag  AnnotationParameterFlag; // Unused. No flags apply
typedef MemberFlag  AliasMemberFlag;      // Unused. No flags apply
typedef MemberFlag  BitflagFlag;         // Unused. No flags apply
typedef MemberFlag  BitsetMemberFlag;    // Unused. No flags apply

// Mask used to remove the flags that do no affect assignability
// Selects T1, T2, O, M, K, D
const unsigned short MemberFlagMinimalMask = 0x003f;

// Flags that apply to type declarationa and DO affect assignability
// Depending on the flag it may not apply to all types
// When not all, the applicable types are listed

```

```

@bit_bound(16)
bitmask TypeFlag {
    @position(0) IS_FINAL,          // F |
    @position(1) IS_APPENDABLE,    // A |- Struct, Union
    @position(2) IS_MUTABLE,       // M | (exactly one flag)

    @position(3) IS_NESTED,        // N   Struct, Union
    @position(4) IS_AUTOID_HASH    // H   Struct
};

typedef TypeFlag StructTypeFlag; // All flags apply
typedef TypeFlag UnionTypeFlag;  // All flags apply
typedef TypeFlag CollectionTypeFlag; // Unused. No flags apply
typedef TypeFlag AnnotationTypeFlag; // Unused. No flags apply
typedef TypeFlag AliasTypeFlag; // Unused. No flags apply
typedef TypeFlag EnumTypeFlag; // Unused. No flags apply
typedef TypeFlag BitmaskTypeFlag; // Unused. No flags apply
typedef TypeFlag BitsetTypeFlag; // Unused. No flags apply

// Mask used to remove the flags that do no affect assignability
const unsigned short TypeFlagMinimalMask = 0x0007; // Selects M, A, F

// Forward declaration
union TypeIdentifier;

// 1 Byte
@extensibility(FINAL) @nested
struct StringTypeDefn {
    SBound          bound;
};

// 4 Bytes
@extensibility(FINAL) @nested
struct StringLTypeDefn {
    LBound          bound;
};

@extensibility(FINAL) @nested

```

```

struct PlainCollectionHeader {
    EquivalenceKind      equiv_kind;
    CollectionElementFlag element_flags;
};

@extensibility(FINAL) @nested
struct PlainSequenceSElemDefn {
    PlainCollectionHeader header;
    SBound                bound;
    @external TypeIdentifier element_identifier;
};

@extensibility(FINAL) @nested
struct PlainSequenceLElemDefn {
    PlainCollectionHeader header;
    LBound                bound;
    @external TypeIdentifier element_identifier;
};

@extensibility(FINAL) @nested
struct PlainArraySElemDefn {
    PlainCollectionHeader header;
    SBoundSeq            array_bound_seq;
    @external TypeIdentifier element_identifier;
};

@extensibility(FINAL) @nested
struct PlainArrayLElemDefn {
    PlainCollectionHeader header;
    LBoundSeq            array_bound_seq;
    @external TypeIdentifier element_identifier;
};

@extensibility(FINAL) @nested
struct PlainMapSTypeDefn {
    PlainCollectionHeader header;
    SBound                bound;
};

```

```

    @external TypeIdentifier element_identifier;
    CollectionElementFlag key_flags;
    @external TypeIdentifier key_identifier;
};

@extensibility(FINAL) @nested
struct PlainMapLTypeDefn {
    PlainCollectionHeader header;
    LBound bound;
    @external TypeIdentifier element_identifier;
    CollectionElementFlag key_flags;
    @external TypeIdentifier key_identifier;
};

// Used for Types that have cyclic dependencies with other types
@extensibility(APPENDABLE) @nested
struct StronglyConnectedComponentId {
    TypeObjectHashId sc_component_id; // Hash StronglyConnectedComponent
    long scc_length; // StronglyConnectedComponent.length
    long scc_index ; // identify type in Strongly Connected Comp.
};

// Future extensibility
@extensibility(MUTABLE) @nested
struct ExtendedTypeDefn {
    // Empty. Available for future extension
};

// The TypeIdentifier uniquely identifies a type (a set of equivalent
// types according to an equivalence relationship: COMPLETE, MNIMAL).
//
// In some cases (primitive types, strings, plain types) the identifier
// is a explicit description of the type.
// In other cases the Identifier is a Hash of the type description
//

```

```

// In the case of primitive types and strings the implied equivalence
// relation is the identity.
//
// For Plain Types and Hash-defined TypeIdentifiers there are three
// possibilities: MINIMAL, COMPLETE, and COMMON:
// - MINIMAL indicates the TypeIdentifier identifies equivalent types
//   according to the MINIMAL equivalence relation
// - COMPLETE indicates the TypeIdentifier identifies equivalent types
//   according to the COMPLETE equivalence relation
// - COMMON indicates the TypeIdentifier identifies equivalent types
//   according to both the MINIMAL and the COMMON equivalence relation.
//   This means the TypeIdentifier is the same for both relationships
//
@extensibility(FINAL) @nested
union TypeIdentifier switch (octet) {
    // ===== Primitive types - use TypeKind =====
    // All primitive types fall here.
    // Commented-out because Unions cannot have cases with no member.
    /*
    case TK_NONE:
    case TK_BOOLEAN:
    case TK_BYTE_TYPE:
    case TK_INT16_TYPE:
    case TK_INT32_TYPE:
    case TK_INT64_TYPE:
    case TK_UINT16_TYPE:
    case TK_UINT32_TYPE:
    case TK_UINT64_TYPE:
    case TK_FLOAT32_TYPE:
    case TK_FLOAT64_TYPE:
    case TK_FLOAT128_TYPE:
    case TK_CHAR8_TYPE:
    case TK_CHAR16_TYPE:
        // No Value
    */

    // ===== Strings - use TypeIdentifierKind =====

```

```

case TI_STRING8_SMALL:
case TI_STRING16_SMALL:
    StringTypeDefn        string_sdefn;

case TI_STRING8_LARGE:
case TI_STRING16_LARGE:
    StringLTypeDefn       string_ldefn;

// ===== Plain collectios - use TypeIdentifierKind =====
case TI_PLAIN_SEQUENCE_SMALL:
    PlainSequenceSElemDefn seq_sdefn;
case TI_PLAIN_SEQUENCE_LARGE:
    PlainSequenceLElemDefn seq_ldefn;

case TI_PLAIN_ARRAY_SMALL:
    PlainArraySElemDefn   array_sdefn;
case TI_PLAIN_ARRAY_LARGE:
    PlainArrayLElemDefn   array_ldefn;

case TI_PLAIN_MAP_SMALL:
    PlainMapSTypeDefn     map_sdefn;
case TI_PLAIN_MAP_LARGE:
    PlainMapLTypeDefn     map_ldefn;

// ===== Types that are mutually dependent on each other ===
case TI_STRONGLY_CONNECTED_COMPONENT:
    StronglyConnectedComponentId sc_component_id;

// ===== The remaining cases - use EquivalenceKind =====
case EK_COMPLETE:
case EK_MINIMAL:
    EquivalenceHash       equivalence_hash;

// ===== Future extensibility =====
// Future extensions
default:
    ExtendedTypeDefn      extended_defn;

```

```

};
typedef sequence<TypeIdentifier> TypeIdentifierSeq;

// --- Annotation usage: -----

// ID of a type member
typedef unsigned long MemberId;
const unsigned long ANNOTATION_STR_VALUE_MAX_LEN = 128;
const unsigned long ANNOTATION_OCTETSEC_VALUE_MAX_LEN = 128;

@extensibility(MUTABLE) @nested
struct ExtendedAnnotationParameterValue {
    // Empty. Available for future extension
};

/* Literal value of an annotation member: either the default value in its
 * definition or the value applied in its usage.
 */

@extensibility(FINAL) @nested
union AnnotationParameterValue switch (octet) {
    case TK_BOOLEAN:
        boolean          boolean_value;
    case TK_BYTE:
        octet            byte_value;
    case TK_INT16:
        short            int16_value;
    case TK_UINT16:
        unsigned short   uint_16_value;
    case TK_INT32:
        long             int32_value;
    case TK_UINT32:
        unsigned long    uint32_value;
    case TK_INT64:
        long long        int64_value;
    case TK_UINT64:
        unsigned long long uint64_value;
};

```



```

    case TK_FLOAT32:
        float                float32_value;
    case TK_FLOAT64:
        double               float64_value;
    case TK_FLOAT128:
        long double         float128_value;
    case TK_CHAR8:
        char                char_value;
    case TK_CHAR16:
        wchar               wchar_value;
    case TK_ENUM:
        long                enumerated_value;
    case TK_STRING8:
        string<ANNOTATION_STR_VALUE_MAX_LEN> string8_value;
    case TK_STRING16:
        wstring<ANNOTATION_STR_VALUE_MAX_LEN> string16_value;
    default:
        ExtendedAnnotationParameterValue extended_value;
};

// The application of an annotation to some type or type member
@extensibility(APPENDABLE) @nested
struct AppliedAnnotationParameter {
    NameHash                paramname_hash;
    AnnotationParameterValue value;
};
// Sorted by AppliedAnnotationParameter.paramname_hash
typedef
sequence<AppliedAnnotationParameter> AppliedAnnotationParameterSeq;

@extensibility(APPENDABLE) @nested
struct AppliedAnnotation {
    TypeIdentifier          annotation_typeid;
    @optional AppliedAnnotationParameterSeq param_seq;
};
// Sorted by AppliedAnnotation.annotation_typeid
typedef sequence<AppliedAnnotation> AppliedAnnotationSeq;

```

```

// @verbatim(placement="<placement>", language="<lang>", text="<text>")
@extensibility(FINAL) @nested
struct AppliedVerbatimAnnotation {
    string<32> placement;
    string<32> language;
    string      text;
};

// --- Aggregate types: -----
@extensibility(APPENDABLE) @nested
struct AppliedBuiltinMemberAnnotations {
    @optional string          unit; // @unit("<unit>")
    @optional AnnotationParameterValue min; // @min , @range
    @optional AnnotationParameterValue max; // @max , @range
    @optional string          hash_id; // @hash_id("<membername>")
};

@extensibility(FINAL) @nested
struct CommonStructMember {
    MemberId                member_id;
    StructMemberFlag        member_flags;
    TypeIdentifier           member_type_id;
};

// COMPLETE Details for a member of an aggregate type
@extensibility(FINAL) @nested
struct CompleteMemberDetail {
    MemberName              name;
    @optional AppliedBuiltinMemberAnnotations ann_builtin;
    @optional AppliedAnnotationSeq          ann_custom;
};

// MINIMAL Details for a member of an aggregate type
@extensibility(FINAL) @nested
struct MinimalMemberDetail {

```

```

        NameHash                                name_hash;
};

// Member of an aggregate type
@extensibility(APPENDABLE) @nested
struct CompleteStructMember {
    CommonStructMember                          common;
    CompleteMemberDetail                        detail;
};

// Ordered by the member_index
typedef sequence<CompleteStructMember> CompleteStructMemberSeq;

// Member of an aggregate type
@extensibility(APPENDABLE) @nested
struct MinimalStructMember {
    CommonStructMember                          common;
    MinimalMemberDetail                        detail;
};

// Ordered by common.member_id
typedef sequence<MinimalStructMember> MinimalStructMemberSeq;

@extensibility(APPENDABLE) @nested
struct AppliedBuiltinTypeAnnotations {
    @optional AppliedVerbatimAnnotation verbatim; // @verbatim(...)
};

@extensibility(FINAL) @nested
struct MinimalTypeDetail {
    // Empty. Available for future extension
};

@extensibility(FINAL) @nested
struct CompleteTypeDetail {
    @optional AppliedBuiltinTypeAnnotations ann_builtin;
    @optional AppliedAnnotationSeq          ann_custom;
    QualifiedTypeName                       type_name;
};

```

```

};

@extensibility(APPENDABLE) @nested
struct CompleteStructHeader {
    TypeIdentifier                base_type;
    CompleteTypeDetail            detail;
};

@extensibility(APPENDABLE) @nested
struct MinimalStructHeader {
    TypeIdentifier                base_type;
    MinimalTypeDetail            detail;
};

@extensibility(FINAL) @nested
struct CompleteStructType {
    StructTypeFlag                struct_flags;
    CompleteStructHeader          header;
    CompleteStructMemberSeq       member_seq;
};

@extensibility(FINAL) @nested
struct MinimalStructType {
    StructTypeFlag                struct_flags;
    MinimalStructHeader           header;
    MinimalStructMemberSeq        member_seq;
};

// --- Union: -----

// Case labels that apply to a member of a union type
// Ordered by their values
typedef sequence<long> UnionCaseLabelSeq;

@extensibility(FINAL) @nested
struct CommonUnionMember {
    MemberId                      member_id;
};

```

```

    UnionMemberFlag          member_flags;
    TypeIdentifier           type_id;
    UnionCaseLabelSeq       label_seq;
};

// Member of a union type
@extensibility(APPENDABLE) @nested
struct CompleteUnionMember {
    CommonUnionMember       common;
    CompleteMemberDetail    detail;
};
// Ordered by member_index
typedef sequence<CompleteUnionMember> CompleteUnionMemberSeq;

// Member of a union type
@extensibility(APPENDABLE) @nested
struct MinimalUnionMember {
    CommonUnionMember       common;
    MinimalMemberDetail     detail;
};
// Ordered by MinimalUnionMember.common.member_id
typedef sequence<MinimalUnionMember> MinimalUnionMemberSeq;

@extensibility(FINAL) @nested
struct CommonDiscriminatorMember {
    UnionDiscriminatorFlag  member_flags;
    TypeIdentifier           type_id;
};

// Member of a union type
@extensibility(APPENDABLE) @nested
struct CompleteDiscriminatorMember {
    CommonDiscriminatorMember    common;
    @optional AppliedBuiltinTypeAnnotations  ann_builtin;
    @optional AppliedAnnotationSeq          ann_custom;
};

```

```

// Member of a union type
@extensibility(APPENDABLE) @nested
struct MinimalDiscriminatorMember {
    CommonDiscriminatorMember    common;
};

@extensibility(APPENDABLE) @nested
struct CompleteUnionHeader {
    CompleteTypeDetail           detail;
};

@extensibility(APPENDABLE) @nested
struct MinimalUnionHeader {
    MinimalTypeDetail            detail;
};

@extensibility(FINAL) @nested
struct CompleteUnionType {
    UnionTypeFlag                union_flags;
    CompleteUnionHeader          header;
    CompleteDiscriminatorMember  discriminator;
    CompleteUnionMemberSeq      member_seq;
};

@extensibility(FINAL) @nested
struct MinimalUnionType {
    UnionTypeFlag                union_flags;
    MinimalUnionHeader           header;
    MinimalDiscriminatorMember   discriminator;
    MinimalUnionMemberSeq       member_seq;
};

// --- Annotation: -----
@extensibility(FINAL) @nested
struct CommonAnnotationParameter {
    AnnotationParameterFlag     member_flags;
    TypeIdentifier               member_type_id;
};

```

```

};

// Member of an annotation type
@extensibility(APPENDABLE) @nested
struct CompleteAnnotationParameter {
    CommonAnnotationParameter    common;
    MemberName                    name;
    AnnotationParameterValue      default_value;
};

// Ordered by CompleteAnnotationParameter.name
typedef
sequence<CompleteAnnotationParameter> CompleteAnnotationParameterSeq;

@extensibility(APPENDABLE) @nested
struct MinimalAnnotationParameter {
    CommonAnnotationParameter    common;
    NameHash                      name_hash;
    AnnotationParameterValue      default_value;
};

// Ordered by MinimalAnnotationParameter.name_hash
typedef
sequence<MinimalAnnotationParameter> MinimalAnnotationParameterSeq;

@extensibility(APPENDABLE) @nested
struct CompleteAnnotationHeader {
    QualifiedTypeName             annotation_name;
};

@extensibility(APPENDABLE) @nested
struct MinimalAnnotationHeader {
    // Empty. Available for future extension
};

@extensibility(FINAL) @nested
struct CompleteAnnotationType {
    AnnotationTypeFlag            annotation_flag;
    CompleteAnnotationHeader      header;
};

```

```

        CompleteAnnotationParameterSeq member_seq;
};

@extensibility(FINAL) @nested
struct MinimalAnnotationType {
    AnnotationTypeFlag          annotation_flag;
    MinimalAnnotationHeader      header;
    MinimalAnnotationParameterSeq member_seq;
};

// --- Alias: -----
@extensibility(FINAL) @nested
struct CommonAliasBody {
    AliasMemberFlag             related_flags;
    TypeIdentifier               related_type;
};

@extensibility(APPENDABLE) @nested
struct CompleteAliasBody {
    CommonAliasBody              common;
    @optional AppliedBuiltinMemberAnnotations ann_builtin;
    @optional AppliedAnnotationSeq ann_custom;
};

@extensibility(APPENDABLE) @nested
struct MinimalAliasBody {
    CommonAliasBody              common;
};

@extensibility(APPENDABLE) @nested
struct CompleteAliasHeader {
    CompleteTypeDetail           detail;
};

@extensibility(APPENDABLE) @nested
struct MinimalAliasHeader {

```



```

    // Empty. Available for future extension
};

@extensibility(FINAL) @nested
struct CompleteAliasType {
    AliasTypeFlag        alias_flags;
    CompleteAliasHeader  header;
    CompleteAliasBody    body;
};

@extensibility(FINAL) @nested
struct MinimalAliasType {
    AliasTypeFlag        alias_flags;
    MinimalAliasHeader  header;
    MinimalAliasBody    body;
};

// --- Collections: -----
@extensibility(FINAL) @nested
struct CompleteElementDetail {
    @optional AppliedBuiltinMemberAnnotations  ann_builtin;
    @optional AppliedAnnotationSeq             ann_custom;
};

@extensibility(FINAL) @nested
struct CommonCollectionElement {
    CollectionElementFlag  element_flags;
    TypeIdentifier         type;
};

@extensibility(APPENDABLE) @nested
struct CompleteCollectionElement {
    CommonCollectionElement  common;
    CompleteElementDetail   detail;
};

@extensibility(APPENDABLE) @nested

```

```

struct MinimalCollectionElement {
    CommonCollectionElement    common;
};

@extensibility(FINAL) @nested
struct CommonCollectionHeader {
    LBound                      bound;
};

@extensibility(APPENDABLE) @nested
struct CompleteCollectionHeader {
    CommonCollectionHeader      common;
    @Optional CompleteTypeDetail detail; // not present for anonymous
};

@extensibility(APPENDABLE) @nested
struct MinimalCollectionHeader {
    CommonCollectionHeader      common;
};

// --- Sequence: -----
@extensibility(FINAL) @nested
struct CompleteSequenceType {
    CollectionTypeFlag          collection_flag;
    CompleteCollectionHeader     header;
    CompleteCollectionElement     element;
};

@extensibility(FINAL) @nested
struct MinimalSequenceType {
    CollectionTypeFlag          collection_flag;
    MinimalCollectionHeader      header;
    MinimalCollectionElement     element;
};

// --- Array: -----
@extensibility(FINAL) @nested

```

```

struct CommonArrayHeader {
    LBoundSeq          bound_seq;
};

@extensibility(APPENDABLE) @nested
struct CompleteArrayHeader {
    CommonArrayHeader  common;
    CompleteTypeDetail detail;
};

@extensibility(APPENDABLE) @nested
struct MinimalArrayHeader {
    CommonArrayHeader  common;
};

@extensibility(APPENDABLE) @nested
struct CompleteArrayType {
    CollectionTypeFlag      collection_flag;
    CompleteArrayHeader     header;
    CompleteCollectionElement element;
};

@extensibility(FINAL) @nested
struct MinimalArrayType {
    CollectionTypeFlag      collection_flag;
    MinimalArrayHeader     header;
    MinimalCollectionElement element;
};

// --- Map: -----
@extensibility(FINAL) @nested
struct CompleteMapType {
    CollectionTypeFlag      collection_flag;
    CompleteCollectionHeader header;
    CompleteCollectionElement key;
    CompleteCollectionElement element;
};

```

```

@extensibility(FINAL) @nested
struct MinimalMapType {
    CollectionTypeFlag      collection_flag;
    MinimalCollectionHeader header;
    MinimalCollectionElement key;
    MinimalCollectionElement element;
};

// --- Enumeration: -----
typedef unsigned short BitBound;

// Constant in an enumerated type
@extensibility(APPENDABLE) @nested
struct CommonEnumeratedLiteral {
    long          value;
    EnumeratedLiteralFlag flags;
};

// Constant in an enumerated type
@extensibility(APPENDABLE) @nested
struct CompleteEnumeratedLiteral {
    CommonEnumeratedLiteral common;
    CompleteMemberDetail    detail;
};

// Ordered by EnumeratedLiteral.common.value
typedef sequence<CompleteEnumeratedLiteral> CompleteEnumeratedLiteralSeq;

// Constant in an enumerated type
@extensibility(APPENDABLE) @nested
struct MinimalEnumeratedLiteral {
    CommonEnumeratedLiteral common;
    MinimalMemberDetail    detail;
};

// Ordered by EnumeratedLiteral.common.value
typedef sequence<MinimalEnumeratedLiteral> MinimalEnumeratedLiteralSeq;

```

```

@extensibility(FINAL) @nested
struct CommonEnumeratedHeader {
    BitBound          bit_bound;
};

@extensibility(APPENDABLE) @nested
struct CompleteEnumeratedHeader {
    CommonEnumeratedHeader common;
    CompleteTypeDetail    detail;
};

@extensibility(APPENDABLE) @nested
struct MinimalEnumeratedHeader {
    CommonEnumeratedHeader common;
};

// Enumerated type
@extensibility(FINAL) @nested
struct CompleteEnumeratedType {
    EnumTypeFlag          enum_flags; // unused
    CompleteEnumeratedHeader header;
    CompleteEnumeratedLiteralSeq literal_seq;
};

// Enumerated type
@extensibility(FINAL) @nested
struct MinimalEnumeratedType {
    EnumTypeFlag          enum_flags; // unused
    MinimalEnumeratedHeader header;
    MinimalEnumeratedLiteralSeq literal_seq;
};

// --- Bitmask: -----
// Bit in a bit mask
@extensibility(FINAL) @nested
struct CommonBitflag {
    unsigned short        position;
};

```

```

        BitflagFlag          flags;
};

@extensibility(APPENDABLE) @nested
struct CompleteBitflag {
    CommonBitflag          common;
    CompleteMemberDetail   detail;
};
// Ordered by Bitflag.position
typedef sequence<CompleteBitflag> CompleteBitflagSeq;

@extensibility(APPENDABLE) @nested
struct MinimalBitflag {
    CommonBitflag          common;
    MinimalMemberDetail   detail;
};
// Ordered by Bitflag.position
typedef sequence<MinimalBitflag> MinimalBitflagSeq;

@extensibility(FINAL) @nested
struct CommonBitmaskHeader {
    BitBound              bit_bound;
};

typedef CompleteEnumeratedHeader CompleteBitmaskHeader;

typedef MinimalEnumeratedHeader MinimalBitmaskHeader;

@extensibility(APPENDABLE) @nested
struct CompleteBitmaskType {
    BitmaskTypeFlag       bitmask_flags; // unused
    CompleteBitmaskHeader header;
    CompleteBitflagSeq    flag_seq;
};

@extensibility(APPENDABLE) @nested
struct MinimalBitmaskType {

```

```

    BitmaskTypeFlag          bitmask_flags; // unused
    MinimalBitmaskHeader     header;
    MinimalBitflagSeq        flag_seq;
};

// --- Bitset: -----
@extensibility(FINAL) @nested
struct CommonBitfield {
    unsigned short          position;
    BitsetMemberFlag       flags;
    octet                   bitcount;
    TypeKind                holder_type; // Must be primitive integer type
};

@extensibility(APPENDABLE) @nested
struct CompleteBitfield {
    CommonBitfield          common;
    CompleteMemberDetail    detail;
};

// Ordered by Bitfield.position
typedef sequence<CompleteBitfield> CompleteBitfieldSeq;

@extensibility(APPENDABLE) @nested
struct MinimalBitfield {
    CommonBitfield          common;
    NameHash                name_hash;
};

// Ordered by Bitfield.position
typedef sequence<MinimalBitfield> MinimalBitfieldSeq;

@extensibility(APPENDABLE) @nested
struct CompleteBitsetHeader {
    CompleteTypeDetail      detail;
};

@extensibility(APPENDABLE) @nested
struct MinimalBitsetHeader {

```

```

        // Empty. Available for future extension
};

@extensibility(APPENDABLE) @nested
struct CompleteBitsetType {
    BitsetTypeFlag        bitset_flags; // unused
    CompleteBitsetHeader header;
    CompleteBitfieldSeq  field_seq;
};

@extensibility(APPENDABLE) @nested
struct MinimalBitsetType {
    BitsetTypeFlag        bitset_flags; // unused
    MinimalBitsetHeader header;
    MinimalBitfieldSeq  field_seq;
};

// --- Type Object: -----
// The types associated with each case selection must have extensibility
// kind APPENDABLE or MUTABLE so that they can be extended in the future

@extensibility(MUTABLE) @nested
struct CompleteExtendedType {
    // Empty. Available for future extension
};

@extensibility(FINAL) @nested
union CompleteTypeObject switch (octet) {
    case TK_ALIAS:
        CompleteAliasType    alias_type;
    case TK_ANNOTATION:
        CompleteAnnotationType annotation_type;
    case TK_STRUCTURE:
        CompleteStructType    struct_type;
    case TK_UNION:
        CompleteUnionType      union_type;
    case TK_BITSET:

```



```

        CompleteBitsetType      bitset_type;
case TK_SEQUENCE:
        CompleteSequenceType   sequence_type;
case TK_ARRAY:
        CompleteArrayType       array_type;
case TK_MAP:
        CompleteMapType        map_type;
case TK_ENUM:
        CompleteEnumeratedType enumerated_type;
case TK_BITMASK:
        CompleteBitmaskType    bitmask_type;

// ===== Future extensibility =====
default:
        CompleteExtendedType    extended_type;
};

@extensibility(MUTABLE) @nested
struct MinimalExtendedType {
    // Empty. Available for future extension
};

@extensibility(FINAL) @nested
union MinimalTypeObject switch (octet) {
    case TK_ALIAS:
        MinimalAliasType        alias_type;
    case TK_ANNOTATION:
        MinimalAnnotationType    annotation_type;
    case TK_STRUCTURE:
        MinimalStructType        struct_type;
    case TK_UNION:
        MinimalUnionType         union_type;
    case TK_BITSET:
        MinimalBitsetType         bitset_type;
    case TK_SEQUENCE:
        MinimalSequenceType      sequence_type;
};

```

```

    case TK_ARRAY:
        MinimalArrayType      array_type;
    case TK_MAP:
        MinimalMapType        map_type;
    case TK_ENUM:
        MinimalEnumeratedType enumerated_type;
    case TK_BITMASK:
        MinimalBitmaskType    bitmask_type;

    // ===== Future extensibility =====
    default:
        MinimalExtendedType    extended_type;
};

@extensibility(APPENDABLE) @nested
union TypeObject switch (octet) { // EquivalenceKind
    case EK_COMPLETE:
        CompleteTypeObject    complete;
    case EK_MINIMAL:
        MinimalTypeObject      minimal;
};
typedef sequence<TypeObject> TypeObjectSeq;

// Set of TypeObjects representing a strong component: Equivalence class
// for the Strong Connectivity relationship (mutual reachability between
// types).
// Ordered by fully qualified typename lexicographic order
typedef TypeObjectSeq      StronglyConnectedComponent;

@extensibility(FINAL) @nested
struct TypeIdentifierTypeObjectPair {
    TypeIdentifier    type_identifier;
    TypeObject        type_object;
};
typedef
sequence<TypeIdentifierTypeObjectPair> TypeIdentifierTypeObjectPairSeq;

```

```

@extensibility(FINAL) @nested
struct TypeIdentifierPair {
    TypeIdentifier type_identifier1;
    TypeIdentifier type_identifier2;
};
typedef sequence<TypeIdentifierPair> TypeIdentifierPairSeq;

@extensibility(APPENDABLE) @nested
struct TypeIdentifierWithSize {
    DDS::Xtypes::TypeIdentifier type_id;
    unsigned long typeobject_serialized_size;
};
typedef sequence<TypeIdentifierWithSize> TypeIdentifierWithSizeSeq;

@extensibility(APPENDABLE) @nested
struct TypeIdentifierWithDependencies {
    TypeIdentifierWithSize typeid_with_size;
    // The total additional types related to minimal_type
    long dependent_typeid_count;
    sequence<TypeIdentifierWithSize> dependent_typeids;
};
typedef
sequence<TypeIdentifierWithDependencies>
TypeIdentifierWithDependenciesSeq;

// This appears in the builtin DDS topics PublicationBuiltinTopicData
// and SubscriptionBuiltinTopicData

@extensibility(MUTABLE) @nested
struct TypeInformation {
    @id(0x1001) TypeIdentifierWithDependencies minimal;
    @id(0x1002) TypeIdentifierWithDependencies complete;
};
typedef sequence<TypeInformation> TypeInformationSeq;

}; // end of module XTypes
}; // end module DDS

```


Annex C: Dynamic Language Binding

The following IDL comprises the API for the Dynamic Language Binding.

```
module DDS {
    local interface DynamicType;
    local interface DynamicTypeBuilder;
    valuetype TypeDescriptor;

    typedef sequence<string> IncludePathSeq;

    local interface DynamicTypeBuilderFactory {
        /*static*/ DynamicTypeBuilderFactory get_instance();
        /*static*/ DDS::ReturnCode_t delete_instance();

        DynamicType get_primitive_type(in TypeKind kind);
        DynamicTypeBuilder create_type(in TypeDescriptor descriptor);
        DynamicTypeBuilder create_type_copy(in DynamicType type);
        DynamicTypeBuilder create_type_w_type_object(
            in TypeObject type_object);
        DynamicTypeBuilder create_string_type(in unsigned long bound);
        DynamicTypeBuilder create_wstring_type(in unsigned long bound);
        DynamicTypeBuilder create_sequence_type(
            in DynamicType element_type,
            in unsigned long bound);
        DynamicTypeBuilder create_array_type(
            in DynamicType element_type,
            in BoundSeq bound);
        DynamicTypeBuilder create_map_type(
            in DynamicType key_element_type,
            in DynamicType element_type,
            in unsigned long bound);
        DynamicTypeBuilder create_bitmask_type(in unsigned long bound);
        DynamicTypeBuilder create_type_w_uri(
            in string document_url,
            in string type_name,
            in IncludePathSeq include_paths);
        DynamicTypeBuilder create_type_w_document(
```

```

        in string document,
        in string type_name,
        in IncludePathSeq include_paths);
    DDS::ReturnCode_t delete_type(in DynamicType type);
};

interface TypeSupport {
// ReturnCode_t register_type(
//     in DomainParticipant domain,
//     in string type_name);
// string get_type_name();

// DynamicType get_type();
};

/* Implied IDL for type "Foo":
interface FooTypeSupport : DDS::TypeSupport {
    DDS::ReturnCode_t register_type(
        in DDS::DomainParticipant participant,
        in string type_name);
    string get_type_name();

    DynamicType get_type();

    Foo create_sample(in DynamicData src);
    DynamicData create_dynamic_sample(in Foo src);
};
*/

interface DynamicTypeSupport : TypeSupport {
    /* This interface shall instantiate the type FooTypeSupport
    * defined by the DDS specification where "Foo" is DynamicData.
    */

    /*static*/ DynamicTypeSupport create_type_support(
        in DynamicType type);
    /*static*/ DDS::ReturnCode_t delete_type_support(

```

```

        in DynamicTypeSupport type_support);

DDS::ReturnCode_t register_type(
    in DDS::DomainParticipant participant,
    in ObjectName type_name);
ObjectName get_type_name();
};

typedef map<ObjectName, ObjectName> Parameters;

valuetype AnnotationDescriptor {
    public DynamicType type;

    DDS::ReturnCode_t get_value(
        inout ObjectName value, in ObjectName key);
    DDS::ReturnCode_t get_all_value(
        inout Parameters value);
    DDS::ReturnCode_t set_value(
        in ObjectName key, in ObjectName value);

    DDS::ReturnCode_t copy_from(in AnnotationDescriptor other);
    boolean equals(in AnnotationDescriptor other);
    boolean is_consistent();
};

valuetype TypeDescriptor {
    public TypeKind kind;
    public ObjectName name;
    public DynamicType base_type;
    public DynamicType discriminator_type;
    public BoundSeq bound;
    @optional public DynamicType element_type;
    @optional public DynamicType key_element_type;

    DDS::ReturnCode_t copy_from(in TypeDescriptor other);
    boolean equals(in TypeDescriptor other);
    boolean is_consistent();
};

```

```

};

valuetype MemberDescriptor {
    public ObjectName name;
    public MemberId id;
    public DynamicType type;
    public string default_value;
    public unsigned long index;
    public UnionCaseLabelSeq label;
    public boolean default_label;

    DDS::ReturnCode_t copy_from(in MemberDescriptor descriptor);
    boolean equals(in MemberDescriptor descriptor);
    boolean is_consistent();
};

local interface DynamicTypeMember {
    DDS::ReturnCode_t get_descriptor(
        inout MemberDescriptor descriptor);

    unsigned long get_annotation_count();
    DDS::ReturnCode_t get_annotation(
        inout AnnotationDescriptor descriptor,
        in unsigned long idx);

    boolean equals(in DynamicTypeMember other);

    MemberId get_id();
    ObjectName get_name();
};

typedef map<ObjectName, DynamicTypeMember> DynamicTypeMembersByName;
typedef map<MemberId, DynamicTypeMember> DynamicTypeMembersById;

local interface DynamicTypeBuilder {
    DDS::ReturnCode_t get_descriptor(
        inout TypeDescriptor descriptor);
};

```



```

ObjectName get_name();
TypeKind get_kind();

DDS::ReturnCode_t get_member_by_name(
    inout DynamicTypeMember member,
    in ObjectName name);
DDS::ReturnCode_t get_all_members_by_name(
    inout DynamicTypeMembersByName member);

DDS::ReturnCode_t get_member(
    inout DynamicTypeMember member,
    in MemberId id);
DDS::ReturnCode_t get_all_members(
    inout DynamicTypeMembersById member);

unsigned long get_annotation_count();
DDS::ReturnCode_t get_annotation(
    inout AnnotationDescriptor descriptor,
    in unsigned long idx);

boolean equals(in DynamicType other);
DDS::ReturnCode_t add_member(in MemberDescriptor descriptor);
DDS::ReturnCode_t apply_annotation(
    in AnnotationDescriptor descriptor);

DynamicType build();
};

local interface DynamicType {
    DDS::ReturnCode_t get_descriptor(
        inout TypeDescriptor descriptor);

    ObjectName get_name();
    TypeKind get_kind();

    DDS::ReturnCode_t get_member_by_name(

```

```

        inout DynamicTypeMember member,
        in ObjectName name);
DDS::ReturnCode_t get_all_members_by_name(
    inout DynamicTypeMembersByName member);

DDS::ReturnCode_t get_member(
    inout DynamicTypeMember member,
    in MemberId id);
DDS::ReturnCode_t get_all_members(
    inout DynamicTypeMembersById member);

unsigned long get_annotation_count();
DDS::ReturnCode_t get_annotation(
    inout AnnotationDescriptor descriptor,
    in unsigned long idx);

    boolean equals(in DynamicType other);
};

local interface DynamicData;

local interface DynamicDataFactory {
    /*static*/ DynamicDataFactory get_instance();
    /*static*/ DDS::ReturnCode_t delete_instance();

    DynamicData create_data();
    DDS::ReturnCode_t delete_data(in DynamicData data);
};

typedef sequence<long>                Int32Seq;
typedef sequence<unsigned long>       UInt32Seq;
typedef sequence<short>               Int16Seq;
typedef sequence<unsigned short>      UInt16Seq;
typedef sequence<long long>           Int64Seq;
typedef sequence<unsigned long long>  UInt64Seq;
typedef sequence<float>               Float32Seq;
typedef sequence<double>              Float64Seq;

```

```

typedef sequence<long double>          Float128Seq;
typedef sequence<char>                 CharSeq;
typedef sequence<wchar>                WcharSeq;
typedef sequence<boolean>              BooleanSeq;
typedef sequence<octet>                ByteSeq;

// typedef sequence<string>            StringSeq;
typedef sequence<wstring>              WstringSeq;

local interface DynamicData {
    readonly attribute DynamicType type;

    DDS::ReturnCode_t get_descriptor(
        inout MemberDescriptor value,
        in MemberId id);
    DDS::ReturnCode_t set_descriptor(
        in MemberId id,
        in MemberDescriptor value);

    boolean equals(in DynamicData other);

    MemberId get_member_id_by_name(in ObjectName name);
    MemberId get_member_id_at_index(in unsigned long index);

    unsigned long get_item_count();

    DDS::ReturnCode_t clear_all_values();
    DDS::ReturnCode_t clear_nonkey_values();
    DDS::ReturnCode_t clear_value(in MemberId id);

    DynamicData loan_value(in MemberId id);
    DDS::ReturnCode_t return_loaned_value(in DynamicData value);

    DynamicData clone();

    DDS::ReturnCode_t get_int32_value(
        inout long value,

```

```

        in MemberId id);
DDS::ReturnCode_t set_int32_value(
    in MemberId id,
    in long value);
DDS::ReturnCode_t get_uint32_value(
    inout unsigned long value,
    in MemberId id);
DDS::ReturnCode_t set_uint32_value(
    in MemberId id,
    in unsigned long value);
DDS::ReturnCode_t get_int16_value(
    inout short value,
    in MemberId id);
DDS::ReturnCode_t set_int16_value(
    in MemberId id,
    in short value);
DDS::ReturnCode_t get_uint16_value(
    inout unsigned short value,
    in MemberId id);
DDS::ReturnCode_t set_uint16_value(
    in MemberId id,
    in unsigned short value);
DDS::ReturnCode_t get_int64_value(
    inout long long value,
    in MemberId id);
DDS::ReturnCode_t set_int64_value(
    in MemberId id,
    in long long value);
DDS::ReturnCode_t get_uint64_value(
    inout unsigned long long value,
    in MemberId id);
DDS::ReturnCode_t set_uint64_value(
    in MemberId id,
    in unsigned long long value);
DDS::ReturnCode_t get_float32_value(
    inout float value,
    in MemberId id);

```

```

DDS::ReturnCode_t set_float32_value(
    in MemberId id,
    in float value);
DDS::ReturnCode_t get_float64_value(
    inout double value,
    in MemberId id);
DDS::ReturnCode_t set_float64_value(
    in MemberId id,
    in double value);
DDS::ReturnCode_t get_float128_value(
    inout long double value,
    in MemberId id);
DDS::ReturnCode_t set_float128_value(
    in MemberId id,
    in long double value);
DDS::ReturnCode_t get_char8_value(
    inout char value,
    in MemberId id);
DDS::ReturnCode_t set_char8_value(
    in MemberId id,
    in char value);
DDS::ReturnCode_t get_char16_value(
    inout wchar value,
    in MemberId id);
DDS::ReturnCode_t set_char16_value(
    in MemberId id,
    in wchar value);
DDS::ReturnCode_t get_byte_value(
    inout octet value,
    in MemberId id);
DDS::ReturnCode_t set_byte_value(
    in MemberId id,
    in octet value);
DDS::ReturnCode_t get_boolean_value(
    inout boolean value,
    in MemberId id);
DDS::ReturnCode_t set_boolean_value(

```

```

        in MemberId id,
        in boolean value);
DDS::ReturnCode_t get_string_value(
    inout string value,
    in MemberId id);
DDS::ReturnCode_t set_string_value(
    in MemberId id,
    in string value);
DDS::ReturnCode_t get_wstring_value(
    inout wstring value,
    in MemberId id);
DDS::ReturnCode_t set_wstring_value(
    in MemberId id,
    in wstring value);

DDS::ReturnCode_t get_complex_value(
    inout DynamicData value,
    in MemberId id);
DDS::ReturnCode_t set_complex_value(
    in MemberId id,
    in DynamicData value);

DDS::ReturnCode_t get_int32_values(
    inout Int32Seq value,
    in MemberId id);
DDS::ReturnCode_t set_int32_values(
    in MemberId id,
    in Int32Seq value);
DDS::ReturnCode_t get_uint32_values(
    inout UInt32Seq value,
    in MemberId id);
DDS::ReturnCode_t set_uint32_values(
    in MemberId id,
    in UInt32Seq value);
DDS::ReturnCode_t get_int16_values(
    inout Int16Seq value,
    in MemberId id);

```

```

DDS::ReturnCode_t set_int16_values(
    in MemberId id,
    in Int16Seq value);
DDS::ReturnCode_t get_uint16_values(
    inout UInt16Seq value,
    in MemberId id);
DDS::ReturnCode_t set_uint16_values(
    in MemberId id,
    in UInt16Seq value);
DDS::ReturnCode_t get_int64_values(
    inout Int64Seq value,
    in MemberId id);
DDS::ReturnCode_t set_int64_values(
    in MemberId id,
    in Int64Seq value);
DDS::ReturnCode_t get_uint64_values(
    inout UInt64Seq value,
    in MemberId id);
DDS::ReturnCode_t set_uint64_values(
    in MemberId id,
    in UInt64Seq value);
DDS::ReturnCode_t get_float32_values(
    inout Float32Seq value,
    in MemberId id);
DDS::ReturnCode_t set_float32_values(
    in MemberId id,
    in Float32Seq value);
DDS::ReturnCode_t get_float64_values(
    inout Float64Seq value,
    in MemberId id);
DDS::ReturnCode_t set_float64_values(
    in MemberId id,
    in Float64Seq value);
DDS::ReturnCode_t get_float128_values(
    inout Float128Seq value,
    in MemberId id);
DDS::ReturnCode_t set_float128_values(

```

```

        in MemberId id,
        in Float128Seq value);
DDS::ReturnCode_t get_char8_values(
    inout CharSeq value,
    in MemberId id);
DDS::ReturnCode_t set_char8_values(
    in MemberId id,
    in CharSeq value);
DDS::ReturnCode_t get_char16_values(
    inout WcharSeq value,
    in MemberId id);
DDS::ReturnCode_t set_char16_values(
    in MemberId id,
    in WcharSeq value);
DDS::ReturnCode_t get_byte_values(
    inout ByteSeq value,
    in MemberId id);
DDS::ReturnCode_t set_byte_values(
    in MemberId id,
    in ByteSeq value);
DDS::ReturnCode_t get_boolean_values(
    inout BooleanSeq value,
    in MemberId id);
DDS::ReturnCode_t set_boolean_values(
    in MemberId id,
    in BooleanSeq value);
DDS::ReturnCode_t get_string_values(
    inout StringSeq value,
    in MemberId id);
DDS::ReturnCode_t set_string_values(
    in MemberId id,
    in StringSeq value);
DDS::ReturnCode_t get_wstring_values(
    inout WstringSeq value,
    in MemberId id);
DDS::ReturnCode_t set_wstring_values(
    in MemberId id,

```



```
        in WstringSeq value);  
}; // local interface DynamicData  
}; // end module DDS
```


Annex D: DDS Built-in Topic Data Types

Previously, the standard DDS type system (based solely on IDL prior to the extensions introduced by this specification) was insufficiently rich to represent the built-in topic data to the level specified by DDS [DDS] and RTPS [RTPS]. This specification remedies this situation. The following are expanded definitions of the built-in topic data types that contain all of the meta-data necessary to represent them as defined by the existing DDS and RTPS specifications.

```
/* dds-xtypes_discovery.idl */

// The types in this file shall be serialized with XCDR encoding version 1
module DDS {
    @extensibility(APPENDABLE) @nested
    struct BuiltinTopicKey_t {
        octet value[16];
    };

    @extensibility(FINAL) @nested
    struct Duration_t {
        long sec;
        unsigned long nanosec;
    };

    @extensibility(APPENDABLE) @nested
    struct DeadlineQosPolicy {
        Duration_t period;
    };

    enum DestinationOrderQosPolicyKind {
        BY_RECEPTION_TIMESTAMP_DESTINATIONORDER_QOS,
        BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS
    };

    @extensibility(APPENDABLE) @nested
    struct DestinationOrderQosPolicy {
        DestinationOrderQosPolicyKind kind;
    };
};
```

```

enum DurabilityQosPolicyKind {
    VOLATILE_DURABILITY_QOS,
    TRANSIENT_LOCAL_DURABILITY_QOS,
    TRANSIENT_DURABILITY_QOS,
    PERSISTENT_DURABILITY_QOS
};

@extensibility(APPENDABLE) @nested
struct DurabilityQosPolicy {
    DurabilityQosPolicyKind kind;
};

enum HistoryQosPolicyKind {
    KEEP_LAST_HISTORY_QOS,
    KEEP_ALL_HISTORY_QOS
};

@extensibility(APPENDABLE) @nested
struct HistoryQosPolicy {
    HistoryQosPolicyKind kind;
    long depth;
};

@extensibility(APPENDABLE) @nested
struct DurabilityServiceQosPolicy {
    Duration_t service_cleanup_delay;
    HistoryQosPolicyKind history_kind;
    long history_depth;
    long max_samples;
    long max_instances;
    long max_samples_per_instance;
};

@extensibility(APPENDABLE) @nested
struct GroupDataQosPolicy {
    ByteSeq value;
};

```

```

};

@extensibility(APPENDABLE) @nested
struct LatencyBudgetQosPolicy {
    Duration_t duration;
};

@extensibility(APPENDABLE) @nested
struct LifespanQosPolicy {
    Duration_t duration;
};

enum LivelinessQosPolicyKind {
    AUTOMATIC_LIVELINESS_QOS,
    MANUAL_BY_PARTICIPANT_LIVELINESS_QOS,
    MANUAL_BY_TOPIC_LIVELINESS_QOS
};

@extensibility(APPENDABLE) @nested
struct LivelinessQosPolicy {
    LivelinessQosPolicyKind kind;
    Duration_t lease_duration;
};

enum OwnershipQosPolicyKind {
    SHARED_OWNERSHIP_QOS,
    EXCLUSIVE_OWNERSHIP_QOS
};

@extensibility(APPENDABLE) @nested
struct OwnershipQosPolicy {
    OwnershipQosPolicyKind kind;
};

@extensibility(APPENDABLE) @nested
struct OwnershipStrengthQosPolicy {
    long value;
};

```

```

};

@extensibility(APPENDABLE) @nested
struct PartitionQosPolicy {
    StringSeq name;
};

enum PresentationQosPolicyAccessScopeKind {
    INSTANCE_PRESENTATION_QOS,
    TOPIC_PRESENTATION_QOS,
    GROUP_PRESENTATION_QOS
};

@extensibility(APPENDABLE) @nested
struct PresentationQosPolicy {
    PresentationQosPolicyAccessScopeKind access_scope;
    boolean coherent_access;
    boolean ordered_access;
};

enum ReliabilityQosPolicyKind {
    BEST_EFFORT_RELIABILITY_QOS,
    RELIABLE_RELIABILITY_QOS
};

@extensibility(APPENDABLE) @nested
struct ReliabilityQosPolicy {
    ReliabilityQosPolicyKind kind;
    Duration_t max_blocking_time;
};

@extensibility(APPENDABLE) @nested
struct ResourceLimitsQosPolicy {
    long max_samples;
    long max_instances;
    long max_samples_per_instance;
};

```

```

@extensibility(APPENDABLE) @nested
struct TimeBasedFilterQosPolicy {
    Duration_t minimum_separation;
};

@extensibility(APPENDABLE) @nested
struct TopicDataQosPolicy {
    ByteSeq value;
};

@extensibility(APPENDABLE) @nested
struct TransportPriorityQosPolicy {
    long value;
};

@extensibility(APPENDABLE) @nested
struct UserDataQosPolicy {
    ByteSeq value;
};

@extensibility(MUTABLE)
struct ParticipantBuiltinTopicData {
    @id(0x0050) @key BuiltinTopicKey_t key;
    @id(0x002C) UserDataQosPolicy user_data;
};

typedef short DataRepresentationId_t;

const DataRepresentationId_t XCDR_DATA_REPRESENTATION = 0;
const DataRepresentationId_t XML_DATA_REPRESENTATION = 1;
const DataRepresentationId_t XCDR2_DATA_REPRESENTATION = 2;

typedef sequence<DataRepresentationId_t> DataRepresentationIdSeq;

const QosPolicyId_t DATA_REPRESENTATION_QOS_POLICY_ID = 23;
const string DATA_REPRESENTATION_QOS_POLICY_NAME = "DataRepresentation";

```

```

@extensibility(APPENDABLE) @nested
struct DataRepresentationQosPolicy {
    DataRepresentationIdSeq value;
};

@bit_bound(16)
enum TypeConsistencyKind {
    DISALLOW_TYPE_COERCION,
    ALLOW_TYPE_COERCION
};

const QosPolicyId_t TYPE_CONSISTENCY_ENFORCEMENT_QOS_POLICY_ID = 24;
const string TYPE_CONSISTENCY_ENFORCEMENT_QOS_POLICY_NAME =
    "TypeConsistencyEnforcement";

@extensibility(APPENDABLE) @nested
struct TypeConsistencyEnforcementQosPolicy {
    TypeConsistencyKind kind;
    boolean ignore_sequence_bounds;
    boolean ignore_string_bounds;
    boolean ignore_member_names;
    boolean prevent_type_widening;
    boolean force_type_validation;
};

@extensibility(MUTABLE)
struct TopicBuiltinTopicData {
    @id(0x005A) @key BuiltinTopicKey_t key;
    @id(0x0005)      ObjectName name;
    @id(0x0007)      ObjectName type_name;
    @id(0x0069) @optional TypeIdV1 type_id; // XTYPES 1.1
    @id(0x0072) @optional TypeObjectV1 type; // XTYPES 1.1
    @id(0x0075) @optional XTypes::TypeInfo type_information;
                                     // XTYPES 1.2
    @id(0x001D)      DurabilityQosPolicy durability;
    @id(0x001E)      DurabilityServiceQosPolicy durability_service;
    @id(0x0023)      DeadlineQosPolicy deadline;
};

```



```

    @id(0x0027)      LatencyBudgetQosPolicy latency_budget;
    @id(0x001B)      LivelinessQosPolicy liveliness;
    @id(0x001A)      ReliabilityQosPolicy reliability;
    @id(0x0049)      TransportPriorityQosPolicy transport_priority;
    @id(0x002B)      LifespanQosPolicy lifespan;
    @id(0x0025)      DestinationOrderQosPolicy destination_order;
    @id(0x0040)      HistoryQosPolicy history;
    @id(0x0041)      ResourceLimitsQosPolicy resource_limits;
    @id(0x001F)      OwnershipQosPolicy ownership;
    @id(0x002E)      TopicDataQosPolicy topic_data;
    @id(0x0073)      DataRepresentationQosPolicy representation;
};

```

```

@extensibility(MUTABLE)
struct TopicQos {
    // ...
    DataRepresentationQosPolicy representation;
};

```

```

@extensibility(MUTABLE)
struct PublicationBuiltinTopicData {
    @id(0x005A) @key BuiltinTopicKey_t key;
    @id(0x0050)      BuiltinTopicKey_t participant_key;
    @id(0x0005)      ObjectName topic_name;
    @id(0x0007)      ObjectName type_name;
    @id(0x0069) @optional TypeIdV1 type_id; // XTYPES 1.1
    @id(0x0072) @optional TypeObjectV1 type; // XTYPES 1.1
    @id(0x0075) @optional XTypes::TypeInfo type_information;
                                     // XTYPES 1.2

    @id(0x001D)      DurabilityQosPolicy durability;
    @id(0x001E)      DurabilityServiceQosPolicy durability_service;
    @id(0x0023)      DeadlineQosPolicy deadline;
    @id(0x0027)      LatencyBudgetQosPolicy latency_budget;
    @id(0x001B)      LivelinessQosPolicy liveliness;
    @id(0x001A)      ReliabilityQosPolicy reliability;
    @id(0x002B)      LifespanQosPolicy lifespan;
    @id(0x002C)      UserDataQosPolicy user_data;
    @id(0x001F)      OwnershipQosPolicy ownership;
};

```

```

    @id(0x0006)      OwnershipStrengthQosPolicy ownership_strength;
    @id(0x0025)      DestinationOrderQosPolicy destination_order;
    @id(0x0021)      PresentationQosPolicy presentation;
    @id(0x0029)      PartitionQosPolicy partition;
    @id(0x002E)      TopicDataQosPolicy topic_data;
    @id(0x002D)      GroupDataQosPolicy group_data;
    @id(0x0073)      DataRepresentationQosPolicy representation;
};

```

```
@extensibility(MUTABLE)
```

```

struct DataWriterQos {
    // ...
    DataRepresentationQosPolicy representation;
};

```

```
@extensibility(MUTABLE)
```

```

struct SubscriptionBuiltinTopicData {
    @id(0x005A) @key BuiltinTopicKey_t key;
    @id(0x0050)      BuiltinTopicKey_t participant_key;
    @id(0x0005)      ObjectName topic_name;
    @id(0x0007)      ObjectName type_name;
    @id(0x0069) @optional TypeIdV1 type_id; // XTYPES 1.1
    @id(0x0072) @optional TypeObjectV1 type; // XTYPES 1.1
    @id(0x0075) @optional XTypes::TypeInformation type_information;
                                     // XTYPES 1.2

    @id(0x001D)      DurabilityQosPolicy durability;
    @id(0x0023)      DeadlineQosPolicy deadline;
    @id(0x0027)      LatencyBudgetQosPolicy latency_budget;
    @id(0x001B)      LivelinessQosPolicy liveliness;
    @id(0x001A)      ReliabilityQosPolicy reliability;
    @id(0x001F)      OwnershipQosPolicy ownership;
    @id(0x0025)      DestinationOrderQosPolicy destination_order;
    @id(0x002C)      UserDataQosPolicy user_data;
    @id(0x0004)      TimeBasedFilterQosPolicy time_based_filter;
    @id(0x0021)      PresentationQosPolicy presentation;
    @id(0x0029)      PartitionQosPolicy partition;
    @id(0x002E)      TopicDataQosPolicy topic_data;
    @id(0x002D)      GroupDataQosPolicy group_data;
};

```

```

        @id(0x0073)      DataRepresentationQosPolicy representation;
        @id(0x0074)      TypeConsistencyEnforcementQosPolicy
                        type_consistency;
};

@extensibility(MUTABLE)
struct DataReaderQos {
    // ...
    DataRepresentationQosPolicy representation;
    TypeConsistencyEnforcementQosPolicy type_consistency;
};
}; // end module DDS

```


Annex E: Built-in Types

DDS shall provide a few very types preregistered “out of the box” to allow users to address certain simple use cases without the need for code generation, dynamic type definition, or type registration. These types are defined below⁹.

```
module DDS {
    @extensibility(APPENDABLE)
    struct _String {
        string value;
    };

    interface StringDataWriter : DataWriter {
        /* This interface shall instantiate the type FooDataWriter defined by
        * the DDS specification where "Foo" is an unbounded string.
        */
    };

    interface StringDataReader : DataReader {
        /* This interface shall instantiate the type FooDataReader defined by
        * the DDS specification where "Foo" is an unbounded string.
        */
    };

    interface StringTypeSupport : TypeSupport {
        /* This interface shall instantiate the type FooTypeSupport
        * defined by the DDS specification where "Foo" is an unbounded
        * string.
        */
    };

    @extensibility(APPENDABLE)
    struct KeyedString {
        @key string key;
        string value;
    };
}
```

⁹ The leading underscore in the declaration of the `String` structure is necessary to prevent collision with the IDL keyword “string.” According to the IDL specification, it is treated as an escaping character and is not considered part of the identifier.

```

};
typedef sequence<KeyedString> KeyedStringSeq;

interface KeyedStringDataWriter : DataWriter {
    /* This interface shall instantiate the type FooDataWriter defined by
     * the DDS specification where "Foo" is KeyedString. It also defines
     * the operations below.
     */
    InstanceHandle_t register_instance_w_key(
        in string key);
    InstanceHandle_t register_instance_w_key_w_timestamp(
        in string key,
        in Time_t source_timestamp);

    ReturnCode_t unregister_instance_w_key(
        in string key);
    ReturnCode_t unregister_instance_w_key_w_timestamp(
        in string key,
        in Time_t source_timestamp);

    ReturnCode_t write_string_w_key(
        in string key,
        in string str,
        in InstanceHandle_t handle);
    ReturnCode_t write_string_w_key_w_timestamp(
        in string key,
        in string str,
        in InstanceHandle_t handle,
        in Time_t source_timestamp);

    ReturnCode_t dispose_w_key(
        in string key);
    ReturnCode_t dispose_w_key_w_timestamp(
        in string key,
        in Time_t source_timestamp);

    ReturnCode_t get_key_value_w_key(

```

```

        inout string key,
        in InstanceHandle_t handle);

InstanceHandle_t lookup_instance_w_key(
    in string key);
};

interface KeyedStringDataReader : DataReader {
    /* This interface shall instantiate the type FooDataReader defined by
    * the DDS specification where "Foo" is KeyedString.
    */
};

interface KeyedStringTypeSupport : TypeSupport {
    /* This interface shall instantiate the type FooTypeSupport
    * defined by the DDS specification where "Foo" is KeyedString.
    */
};

@extensibility(APPENDABLE)
struct Bytes {
    ByteSeq value;
};

typedef sequence<Bytes> BytesSeq;

interface BytesDataWriter : DataWriter {
    /* This interface shall instantiate the type FooDataWriter defined by
    * the DDS specification where "Foo" is an unbounded sequence of
    * bytes (octets). It also defines the operations below.
    */
    ReturnCode_t write_w_bytes(
        in ByteArray bytes,
        in long offset,
        in long length,
        in InstanceHandle_t handle);
    ReturnCode_t write_w_bytes_w_timestamp(

```

```

        in ByteArray bytes,
        in long offset,
        in long length,
        in InstanceHandle_t handle,
        in Time_t source_timestamp);
};

interface BytesDataReader : DataReader {
    /* This interface shall instantiate the type FooDataReader defined by
    * the DDS specification where "Foo" is Bytes.
    */
};

interface BytesTypeSupport : TypeSupport {
    /* This interface shall instantiate the type FooTypeSupport
    * defined by the DDS specification where "Foo" is Bytes.
    */
};

@extensibility(APPENDABLE)
struct KeyedBytes {
    @key string key;
    ByteSeq value;
};
typedef sequence<KeyedBytes> KeyedBytesSeq;

interface KeyedBytesDataWriter : DataWriter {
    /* This interface shall instantiate the type FooDataWriter defined by
    * the DDS specification where "Foo" is KeyedBytes. It also defines
    * It also defines the operations below.
    */
    InstanceHandle_t register_instance_w_key(
        in string key);
    InstanceHandle_t register_instance_w_key_w_timestamp(
        in string key,
        in Time_t source_timestamp);
};

```



```

ReturnCode_t unregister_instance_w_key(
    in string key);
ReturnCode_t unregister_instance_w_key_w_timestamp(
    in string key,
    in Time_t source_timestamp);

ReturnCode_t write_bytes_w_key(
    in string key,
    in ByteArray bytes,
    in long offset,
    in long length,
    in InstanceHandle_t handle);
ReturnCode_t write_bytes_w_key_w_timestamp(
    in string key,
    in ByteArray bytes,
    in long offset,
    in long length,
    in InstanceHandle_t handle,
    in Time_t source_timestamp);

ReturnCode_t dispose_w_key(
    in string key);
ReturnCode_t dispose_w_key_w_timestamp(
    in string key,
    in Time_t source_timestamp);

ReturnCode_t get_key_value_w_key(
    inout string key,
    in InstanceHandle_t handle);

InstanceHandle_t lookup_instance_w_key(
    in string key);
};

interface KeyedBytesDataReader : DataReader {
    /* This interface shall instantiate the type FooDataReader defined by

```

```
    * the DDS specification where "Foo" is KeyedBytes.
    */
};

interface KeyedBytesTypeSupport : TypeSupport {
    /* This interface shall instantiate the type FooTypeSupport
    * defined by the DDS specification where "Foo" is KeyedBytes.
    */
};

}; // end module DDS
```

Annex F: Characterizing Legacy DDS Implementations

Prior to the adoption of this specification, no formal definition existed of the DDS Type System or of those portions of IDL that corresponded to it. This annex provides a non-normative description of what is believed to be the consensus Type System, Type Representation, Data Representation, and Language Binding of DDS implementations that do not conform to this specification. It is provided for the convenience of implementers and evaluators who may wish to compare and contrast DDS implementations or to distinguish those parts of this specification that are novel from those that merely codify previous de-facto-standard practice.

F.1 Type System

The following portions of the Type System are believed to be supported by the majority of DDS implementations, regardless of their compliance with this specification:

- *Namespaces and modules.*
- *All primitive types*, albeit named according to their mappings in the IDL Type Representation.
- *Enumerations* of bit bound 32 with automatically assigned enumerator values.
- *Aliases*, typically referred to as “`typedefs`” based on their mappings in the IDL Type Representation.
- *Arrays*, both single-dimensional and multi-dimensional.
- *Sequences*, both bounded and unbounded.
- *Strings* of narrow or wide characters, both bounded and unbounded.
- *Structures* without inheritance. User-defined structures have `FINAL` extensibility. Members are typically non-optional, non-shared, and do not expose member IDs. DDS-RTPS-compliant implementations support `MUTABLE` extensibility and the `must_understand` attribute with respect to the built-in topic data types. Otherwise, these attributes are not generally supported. Key members are generally supported.
- *Unions* with `FINAL` extensibility and without key members. Discriminators of wide character and octet types are not generally supported.

F.2 Type Representation

The IDL Type Representations of those portions of the Type System enumerated above are generally supported.

The XSD Type Representation is based heavily on the “CORBA to WSDL/SOAP Interworking Specification” and as such may to some extent be said to predate this specification. However, support for representing types in XSD is not widespread among DDS implementations that do not comply with this specification.

F.3 Data Representation

The Extended CDR Representations of those portions of the Type System enumerated above are generally supported. The exception is the extended parameter ID and length facility based on `PID_EXTENDED`, which is not generally supported.

F.4 Language Binding

The Plain Language Bindings of those portions of the Type System enumerated above are generally supported.